MAINVIEW® for IMS Online
Monitors and Traces Reference Manual

Version 3.3.30

November 2003

Contacting BMC Software

You can access the BMC Software Web site at **http://www.bmc.com**. From this Web site, you can obtain information about the company, its products, corporate offices, special events, and career opportunities.

United States and Canada

Address	BMC Software, Inc. 2101 CityWest Blvd. Houston TX 77042-2827
Telephone	713 918 8800 or 800 841 2031
Fax	713 918 8000

Outside United States and Canada

Telephone	(01) 713 918 8800
Fax	(01) 713 918 8000

Customer Support

You can obtain technical support by using the Support page on the BMC Software Web site or by contacting Customer Support by telephone or e-mail. To expedite your inquiry, please see “Before Contacting BMC Software.”

Support Web Site

You can obtain technical support from BMC Software 24 hours a day, 7 days a week at **http://www.bmc.com/support.html**. From this Web site, you can

- read overviews about support services and programs that BMC Software offers
- find the most current information about BMC Software products
- search a database for problems similar to yours and possible solutions
- order or download product documentation
- report a problem or ask a question
- subscribe to receive e-mail notices when new product versions are released
- find worldwide BMC Software support center locations and contact information, including e-mail addresses, fax numbers, and telephone numbers

Support by Telephone or E-mail

In the United States and Canada, if you need technical support and do not have access to the Web, call 800 537 1813. Outside the United States and Canada, please contact your local support center for assistance. To find telephone and e-mail contact information for the BMC Software support center that services your location, refer to the Contact Customer Support section of the Support page on the BMC Software Web site at **http://www.bmc.com/support.html**.

Before Contacting BMC Software

Before you contact BMC Software, have the following information available so that Customer Support can begin working on your problem immediately:

- product information
  - product name
  - product version (release number)
  - license number and password (trial or permanent)
- operating system and environment information
  - machine type
  - operating system type, version, and service pack or other maintenance level such as PUT or PTF
  - system hardware configuration
  - serial numbers
  - related software (database, application, and communication) including type, version, and service pack or maintenance level
- sequence of events leading to the problem
- commands and options that you used
- messages received (and the time and date that you received them)
  - product error messages
  - messages from the operating system, such as `file system full`
  - messages from related software

Contents

Figures

Tables

How to Use This Book

This manual documents the features and functions of the MAINVIEW® for IMS (MVIMS) Online Resource Monitor, Workload Monitor, and trace services.

The features, and functions of the MVIMS Online Resource Analyzer and Workload Analyzer display services are documented in the *MAINVIEW for IMS Online – Analyzers Reference Manual*.

The features, and functions of the IMSPlex System Manager (IPSM) component are documented in the *MAINVIEW for IMS Online – IPSM Reference Manual*.

For information about what's new in the current release of MAINVIEW for IMS Online, see the product Release Notes, which are available on the BMC Software Support Web pages.

This manual is intended for use by the IMS master terminal operator (MTO), system programmer, database administrator, or performance analyst who monitors the status, activity, and performance of IMS and its resources.

MVIMS Product Library

MVIMS is integrated with the BMC Software MAINVIEW® architecture. MAINVIEW is a base architecture that allows authorized users to use a single terminal to interrogate any OS/390, CICS, IMS, DB2, or MQSeries subsystem in a sysplex.

The MVIMS product library includes the following documents.

MAINVIEW for IMS Online:

MAINVIEW for IMS Online – Customization Guide

MAINVIEW for IMS Online – IPSM Reference Manual

MAINVIEW for IMS Online – Analyzers Reference Manual

MAINVIEW for IMS Online – Monitors and Traces Reference Manual

MAINVIEW for IMS Online – Release Notes

MAINVIEW for IMS Offline:

MAINVIEW for IMS Offline – Customization and Utilities Guide

MAINVIEW for IMS Offline – Performance Reporter Reference Manual

MAINVIEW for IMS Offline – Transaction Accountant Reference Manual

MAINVIEW for IMS Offline – Release Notes

How This Book Is Organized

This manual is divided into the following parts:

- Part 1 describes techniques to help you use MVIMS to optimize IMS performance. It contains references to specific services.
- Part 2 describes how to use MVIMS. It includes descriptions of
  - Analyzers and monitors and how they are used
  - The online functions that can be used with MVIMS in a terminal session (TS)
  - The Primary Option Menu and applications that provide easy access to product service applications
- Part 3 describes how to use each of the monitor data collection services.
- Part 4 describes how trace services are used to activate and display a summary workload trace or detail transaction processing event trace.
- Part 5 describes how to control time-driven monitor services.
- Part 6 has appendixes that provide information about how to analyze IMS dumps and how to use the MAINVIEW distributed product libraries.

To help you find information about a service quickly:

- Service descriptions are arranged alphabetically by name and grouped by sections that correspond to a major IMS transaction processing event.
- Service parameters, display fields, and DWAIT display events are described in alphabetical order.
- All of the service names are in alphabetical order together as indexed entries to “Service Select Code” in the index.

Related MAINVIEW Products

The related MAINVIEW-based products include the following:

- MAINVIEW® AutoOPERATOR™
- MAINVIEW® for CICS
- MAINVIEW® for DB2®
- MAINVIEW® for DBCTL
- MAINVIEW® FOCAL POINT
- MAINVIEW® for MQSeries
- MAINVIEW® for OS/390
- MAINVIEW® VistaPoint™

Customization and administration instructions for the MAINVIEW-based functions are provided in the *MAINVIEW Common Customization Guide*. The following manuals document product-specific customization instructions:

- *MAINVIEW AutoOPERATOR Customization Guide*
- *MAINVIEW for CICS Customization Guide*
- *MAINVIEW for DB2 Customization Guide*
- *MAINVIEW for DBCTL Customization Guide*
- *MAINVIEW for IMS Online – Customization Guide*
- *MAINVIEW for IMS Offline – Customization and Utilities Guide*
- *MAINVIEW for OS/390 Customization Guide*

The following books document the use of general services common to MAINVIEW for IMS and related products:

- *MAINVIEW AutoOPERATOR Basic Automation Guide*
- *MAINVIEW AutoOPERATOR Advanced Automation Guide for CLIST EXECS*
- *MAINVIEW AutoOPERATOR Advanced Automation Guide for REXX EXECS*
- *MAINVIEW for CICS PERFORMANCE REPORTER User Guide*
- *MAINVIEW for DB2 User Guide*
- *MAINVIEW for DBCTL Analyzers, Monitors, and Traces Reference Manual*

Related IBM® Publications

OS/390 Initialization and Tuning Guide
IMS Operator Reference
System Administration Guide

Conventions Used in This Manual

The following symbols are used to define command syntax, are ***not*** part of the command, and should never be typed as part of the command:

- Brackets [ ] enclose optional parameters or keywords.
- Braces { } enclose a list of parameters; one must be chosen.
- A line | separates alternative options; one can be chosen.
- An <u>underlined</u> parameter is the default.

The following command syntax conventions apply:

- An ITEM IN CAPITAL LETTERS must be typed exactly as shown.
- Items in *italicized, lowercase* letters are values that you supply.
- When a command is shown in uppercase and lowercase letters, such as **HSplit**, the uppercase letters show the command abbreviation that you can use (**HS**, for example). The lowercase letters complete the entire command name. Typing the entire command name is an optional, alternative way of entering the command.
- Commands without an abbreviation (**END**, for example) appear in all uppercase letters.

Part 1. Performance Analysis and Monitoring Techniques

This section describes techniques to help you use MVIMS to optimize system performance.

Chapter 1. Optimizing System Workflow

This chapter describes how MVIMS can be used to analyze and monitor system operation and performance. It is organized into sections by IMS processing event components: MFS, queuing, scheduling, operations (region activity), database, IMS internal functions, and OS/390 functions. Within each section, a performance problem is presented with the applicable analyzer or monitor solution.

MVIMS Usage Modes

MVIMS can be used in two different modes:

- For ongoing analysis of what is currently happening in IMS.

  This mode provides a realtime picture of the current state of the system. These realtime displays are provided by Resource Analyzer and Workload Analyzer services.

  Analyzer services can be time-driven. In this mode, you can refresh the displays in a time-driven cycle or log them automatically at time-driven intervals to a BBI-SS PAS Image log for later retrieval.

- As a monitor to assist you in tuning the performance of the whole system by investigating the functions and resources over time that may be bottlenecks.

  This mode provides statistics accumulated at user-specified time intervals. Time-driven IMS workload samplings and graphic plot displays of the collected data samplings are provided by Workload Monitor and Resource Monitor services. IMS workload wait event and transaction trace data are collected and displayed by Workload Analyzer.

  Monitor-collected values can be compared to user-specified thresholds and warning messages can be generated. A warning message can invoke automatic corrective action from MAINVIEW AutoOPERATOR or alert a user to take corrective action.

You can set MVIMS services to monitor IMS performance at different times of the day or on different days of the week when processing characteristics may change. Small samples can be taken over time and comparisons made to determine the best performance indicators. Service sampling of IMS performance can help you determine if some changes are caused by application program or system design errors or oversight. Concentration should be directed in the areas where changes have the most effect in system performance.

Often it is valuable to know what has happened in a particular time interval (for example, the last 5-10 minutes). This information can be determined by making two observations and calculating how many events have occurred between the two. By relating these numbers to the length of the time interval chosen, rates such as requests-per-second or I/O-per-second can be calculated. Rates such as I/O-per-request can be found by comparing the number of events to any other counter (possible for a specific time span also). These calculations require more work, but result in more meaningful measurements and numbers that can be compared over time to show changes in the system.

With Resource Monitor, this work is performed automatically for the most important system variables, and the measurement made by each can be displayed with the PLOT service.
If several monitor requests are set up for the same sampling interval, the current measurement values can be viewed easily with the DMON service for comparisons.

In some subsections, suggestions are made for ways to improve performance. These should be considered as a starting point and not as a list of the best, or only, actions to be taken. A suggestion may not be valid in every situation and must be considered for its value in the particular installation and mode of operation.

Message Format Service (MFS)

This section describes analyzing and monitoring MFS.

Analyzing MFS

Message Format Service (MFS) is the first major IMS function encountered by an incoming request (or the last for an outgoing response). MFS can have a great impact upon the efficiency and productivity of the entire system. As one of the unique features of IMS/TM, the generality and flexibility of MFS processing can result in performance problems if misused. For optimum performance, define only as many MFS formats as will fit in the MFS pool.

The formats reside online as records in a PDS with `RECFM=U`. The MFS data set should be allocated as a single extent by cylinders to an area without alternate tracks. The volume should be mounted private on nonshared DASD.

When a format is requested that is not in the pool, MFS must locate the directory entry for the format. If there is an entry for it in the in-core directory $$IMSDIR, a directory read is not necessary to retrieve this entry from the directory block.

This entry contains the TTR of the first record of the format and a half word of user data that contains the total length of the format block. Using this length, space is obtained in the MFS pool to contain the block. The format is then read into this area, one record at a time. The format block may be split into multiple records if the block size is too small. The block size should be at least as large as the largest format block because additional I/O is required to retrieve these multi-record formats.

The formats are maintained in the pool as long as possible to reduce the amount of I/O. Only when space is required in the pool does MFS free the least-recently-referenced unused format block. This release continues until sufficient contiguous free space has been obtained to satisfy the request. Fragmentation can be a problem in any pool of variable length blocks managed in this manner. The service Message Format Pool Utilization (MFSUT) can be used to study the effect of space release upon the MFS pool, including fragmentation. For more information about MFSUT, see the service description in Part 3, “Analyzers,” in *MAINVIEW for IMS Online – Analyzers Reference Manual*.

The basic request types that can be made to MFS are

- PRE-FETCH REQUESTS
- IMMEDIATE-REQUESTS
- FREE BLOCK REQUESTS

PRE-FETCH is an anticipated future need for a format block. If the pre-fetch feature is enabled, it informs PRE-FETCH to retrieve the requested format block.

An IMMEDIATE REQUEST is for a block that must be read into the MFS buffer pool before processing can continue. The IMMEDIATE REQUESTs counter is a very good indicator of MFS activity, especially in relation to total MFS reads, and is useful in evaluating MFS performance. It is possible to calculate the average physical I/O-per-second to the MFS data set. Depending on the device and contention, this statistic may or may not indicate a problem.

FREE BLOCK REQUESTS inform the MFS pool handler that the format block is no longer actively being referenced and is now a candidate for being washed from the pool if space is needed for another block.

The MFSST service displays these counters in matrix format (for more information about MFSST, see the “Message Format Service Displays” chapter in the *MAINVIEW for IMS Online – Analyzers Reference Manual*). Each line represents one of the three request types:

- PRE-FETCH REQUEST
- IMMEDIATE REQUEST
- FREE BLOCK REQUEST

Each column represents one of the three major queues where the FRE could have been located:

- PRE-FETCH QUEUE
- IMMEDIATE QUEUE
- FREE BLOCK QUEUE

Currently, the counter is incremented when the FRE is located in a queue. This does not necessarily mean that the format block itself is already in the pool. The MFS fetch request handler DFSFFRH0 could be modified to check for a loaded format block before incrementing this counter. This can lead to a better understanding of MFS performance. The percent of IMMEDIATE REQUESTs satisfied in the pool (in the IMMEDIATE or FREE BLOCK queues) is a good indicator of how well MFS is performing. The larger the percentage, the fewer synchronous waits for format I/O.

The MFS pool is above the 16Mb line. Because of virtual storage constraint relief (VSCR) above the 16Mb line, the MFS pool can be increased by 400 to 500 percent to further reduce the number of I/Os to the FORMAT data set. If the pool size is increased, be sure to monitor the number of FREs because it may be necessary to increase them to utilize the additional space in the pool. The directory entries are built dynamically in the pool as they are used. The $$IMSDIR is still used, but is not as significant as in previous releases. These dynamic entries are flushed whenever an online change is made to the MFS formats.

Monitoring MFS

MFS performance can be monitored continuously by these Resource Monitor monitors:

MFSIO MFS input/output requests
MFSIR MFS immediate requests
MFSFD Percent of MFS blocks found in pool

If any of these values are considered to be excessive in given time periods, additional investigation is required (see “Suggestions” on page 6). There are no specific values that are good or bad; each system must be evaluated individually.

Suggestions

- Optimize the pool space and FRE allocation as described.
- Use the in-core format directory $$IMSDIR.
- Eliminate any unnecessary or unused format blocks.
- After updating/deleting a format, compress the library.
- Allocate by cylinder (but no more space than necessary) and make certain there are no alternate tracks.
- Watch the placement of the MFS data set and mount private on nonshared DASD.
- Make certain that the block size is at least as large as the largest format.
- Do not allocate any more directory blocks than necessary.
- If the MFS data set is large, consider reordering the formats. You can determine an optimal order with one of these methods:
  - Analyze the IMS log to determine the frequency of use.
  - Analyze the DC Monitor output to calculate the frequency fetched.

  You can then build the MFS data set by generating the formats in order of decreasing use. Use the frequency fetched from the DC Monitor to reorder the individual format blocks.

Queuing

This section discusses analyzing and monitoring queuing.

Analyzing Queuing

The next major IMS function that can be a major system bottleneck is the queuing of input requests and output responses. Like MFS, the queuing routines of IMS attempt to keep as much as possible in the queue pool to reduce I/O. I/O is done only when checkpoints request a pool purge, when space is needed in the pool, or when something previously written out must be retrieved. Thus, the queue pool and data sets can be critical to good performance.

Queuing makes use of preformatted OSAM data sets.

Each data set should be allocated by cylinder as a single extent on an area without any alternate tracks.

- These data sets should be mounted private on a low-contention, nonshared device.

The optimization of the queue pool parameters is a more difficult problem. Several problem areas are:

- There are constraints that limit the possible values that can be defined. IMS imposes a strict relationship between the block size (BLKSIZE) and logical record length (LRECL) of the three queue data sets (space is wasted unless the block size is a multiple of all three LRECLs).

- The minimum LRECL of the LONG message queue data set frequently is dictated by the applications. The length of segments being queued to and from the applications determines the optimal choice of the three parameters.
- The LRECL for the SHORT message queue data set is especially difficult to determine. If this LRECL is too large, the SHORT message queue data set will be overutilized. If the LRECL is too small, the LONG message queue data set will be overutilized. Both situations waste space.

MVIMS can point to a possible problem in this area, but the solution may require offline analysis to determine the average segment length or the segment length distribution. These online Analyzer displays are helpful:

- Service DREGN (PSB/Transaction area of display) shows the average length of all input messages of the transaction types currently processing (as calculated by the IMS queue manager).
- Service STAT shows the utilization of each of the three queue data sets. By observing these percentages for a period of time, an imbalance between SHORT and LONG can be detected easily.

To maximize the use of the queue pool and reduce I/O, messages that remain queued for a long time should be set up as candidates to wash out. This can be done by setting the record length of the LONG message queue data set equal to the common block size and defining the segment length of such messages long enough to force them into the LONG message queue. This allows a block filled by such a message to be written out immediately, releasing the block. Otherwise, the segment takes up only one record in the block and the block may remain in the pool, but only the remaining records can be used.

The types of messages that should be considered are transactions that normally do not schedule (locked, stopped, or priority zero), and long output messages that are routed to slow remote printers or that use terminal operator paging. Any space wasted on the data set by possible padding is of minor importance in comparison to the better use of the queue buffers in the pool.

The remaining parameter is the determination of the number of queue buffers. Too few queue buffers results in unnecessary queue manager I/O. Too many queue buffers waste central storage and increase the paging rate on a real-storage constrained system.

Although the TOTAL REQUESTS to the queue manager is the best indication of the amount of queuing activity, the number of ENQUEUES plus CANCELS gives a better idea of the number of messages involved. These messages can be either single/multiple segment input requests or output responses. Using this sum, the amount of queuing activity per message can be calculated (TOTAL REQUESTS/(ENQUEUES + CANCELS)).

It is also possible to calculate the amount of queue manager activity per second. REPOSITIONS is a nonproductive call to the queue manager used to re-examine a previous segment. The number of PCBs UNCHAINED is incremented for each PCB that loses position when a buffer is written out. Finally, the number of WAITS FOR AN AVAILABLE BUFFER is incremented when a request had a wait while a buffer is written out to release space. As with many statistics, they are more meaningful when viewed as events per second or events per unit of work (such as I/O per message).

Monitoring Queuing

There are 15 Resource Monitor services for monitoring queuing data (see Part 3, "Monitors). Workload Monitor services monitor the input queue time by various selection criteria, such as class, transaction code, LTERM, region, program, and USERID. With well-considered threshold specification, these monitors produce warning messages that can be used to trigger automatic operational changes or alert operations personnel to use analyzer services to investigate further.

Scheduling

This section describes analyzing and monitoring scheduler pool utilization.

Analyzing Scheduling

The scheduling function of IMS (selecting and preparing application programs to run in each dependent region as it completes its previous task) is very complex and critical for good performance of the system. It is dependent on several factors that you can directly influence (for example, pool sizes and transaction class assignments). However, the basis of information for making decisions that have positive results on this process are difficult to obtain when it is needed. MVIMS services such as SCHED, STAT, and DSPST address several of the major problem areas.

Class Queuing

IMS queues and schedules transactions according to class assignments and by priority within class if needed. Each transaction is defined to the system as belonging to one class. Each message region has from one to four specified classes that it can process. This allows you to balance the processing load, give priority to transactions with critical response time requirements, isolate long-running transactions, and so on.

The CLASQ display presents an overview of the current situation: what is queued in each of the classes, what the regions are doing currently, and what classes each region can accept.

By examining the status of the queues at various times of the day, any imbalance caused by the class assignments should be indicated by excessive queue size for some classes or by idle regions waiting for input in other classes. If such imbalances occur often (perhaps only at certain times of the day), performance can be improved either by reassigning transaction classes or by dynamically changing the region processing classes. Because this display shows the status of all the classes at once, the effect of such reassignments can be detected easily and followed through time.

Note: An increasing queue delay caused by load imbalance (check the enqueue time (ENQ TIM) of the current transactions being processed) also impacts the system by increasing the overhead incurred by the queue manager. The longer the queues, the greater the possibility that a transaction will be washed out of the pool to the queue data sets before it is needed, causing extra I/O activity to write it out and to retrieve it. The QUEST service can be used to investigate this occurrence (see "Analyzing Queuing" on page 6). See the information in the REPOSITIONS or PCBs UNCHAINED fields of the QUEST display.

Balancing Group (BALG) Queuing

IMS Fast Path transactions do not use the message queues. Instead of the message queuing mechanism, each physical terminal has an expedited message handler (EMH) buffer associated with it. When an input message is received by IMS and is defined either as Fast Path exclusive or Fast Path potential, control is given to a user exit that can modify the routing code. Assign a different routing code or indicate that the message should be processed as a full-function IMS transaction; that is, use normal message queuing.

Once a Fast Path routing code is assigned to an input message, the balancing group (BALG) is determined and the EMH buffer is queued to the BALG for processing.

This queuing takes place only if the BALG is active. The EMH buffers are processed in a FIFO sequence by the Fast Path regions that service the BALG. More than one region can process the same BALG. By observing the data displayed by the BALGQ service, it is possible to determine when there is a delay in processing for a specific BALG by the presence of a queue. This problem can be remedied by increasing the number of regions processing the BALG.

Scheduling Activity

When a region is free and input belonging to one of its classes is available, IMS attempts to schedule a transaction. Many checks are necessary to ensure successful scheduling, and if any of these checks fail, that region may have to wait for some event before it can become active. However, in this case, the attempt is made first to schedule a different transaction.

Monitoring the success of system scheduling over time (by viewing the percentages of schedules and failures on the Resource Analyzer SCHED display or calculating the number of schedulings per minute) can point out possible problems before they become critical or make reduction of hidden overhead possible.

- PROGRAM CONFLICTS can be reduced by allowing parallel processing (load balancing); however, this can cause PI conflicts because of database intent. Care must be taken to minimize both of these items.
- If PRIORITY CUTOFFS are occurring regularly, check these definitions carefully to see if it is actually necessary to bypass available transactions (this can cause a region to be idle) to wait for the higher priority transaction.
- INTENT FAILURES can be caused either by insufficient space in the PSB, PSB work, or DMB pools, or by database intent failures. Excessive intent failures can indicate that too many regions are competing for the available resources, increasing the time needed for each, both in scheduling and in processing. This could be the result of allowing parallel scheduling (see “Program Conflicts” above).

  Note: Currently, database intent failures occur only when a PROCOPT of EXCLUSIVE is used.
- OTHER REASONS for failures include locked or stopped transactions, programs, or databases.

  Note: The associated SMBs are only removed from scheduling queues if the program is bad (BLDL failed) or a checkpoint purge is in progress. Although locked transactions no longer appear as schedulable in the CLASQ display, the SMBs are still examined and this counter is incremented. To prevent this counter from being continuously incremented, assign such transactions to an unassigned processing class.

Suggestions

- Complicated class and priority assignments or the use of many dependent regions increases the contention during scheduling. The simple way is often the best.
- Use the processing limit for key transactions to reduce the number of schedulings. The same program being rescheduled causes a program fetch each time (unless preloaded), and is very expensive.
- Consider Wait for Input (WFI) processing for frequently used transactions. This dedicates a dependent region(s) to a transaction and eliminates program fetch.

Monitoring Scheduling

Several Resource Monitor monitors are available to monitor arrival rates and processing that can affect scheduling. For example, it might be necessary to stop a less important transaction or BMP if the arrival rate of a particular transaction exceeds a user-defined threshold. This can be done either by an MAINVIEW AutoOPERATOR for IMS EXEC or through manual intervention.

Pool Utilization

This section describes PSB and DMB pools and pool utilization for `LSO=S`.

PSB and DMB Pools

The two main pools concerned in scheduling contain the IMS control blocks that define the application program's logical databases (PSBs) and the physical databases they access (DMBs) for the DL/I (IMS Data Language/I) interface. If not already present, the necessary blocks must be loaded into these pools during scheduling. A pool space failure can cause a region being scheduled to wait until resources are released by the completion of work in another region. Check the intent failure counter in the SCHED display.

If a DMB must be washed from the pool to free space for another, the associated data sets must be closed (a very time-consuming process that should be avoided). If that DMB is ever accessed again, it is necessary to reload it and open the associated data sets. If LSO=Y is coded, all open/close activity occurs in the IMS control region.

The PSBUT and DMBUT displays (see the service descriptions in Part 3, “Analyzers,” in the *MAINVIEW for IMS Online – Analyzers Reference Manual*) are designed to show the free space still available in the pools and fragmentation and its potential impact. The displays show the current status of the pool and simulate the results of applying the least-recently-referenced space release algorithm to the blocks not currently being used. Each line displays successively the number of allocated blocks and the number of free spaces in the pool with minimum, average, maximum, and total lengths.

With these displays, it is possible to see not only the current free space (first line), but also the total free space available if needed (last line). This free space would be gained by freeing allocated but unused blocks. The maximum free space column shows the largest block that could be loaded into the pool (compare this to the maximum defined block size).

Pool Utilization for LSO=S

Open/close activity takes place in the DLISAS region and is still DLISAS time-consuming, but it does not impact IMS work in the control region. If `LSO=S` is selected, the PSB pool is split into two separate pools. One of these pools resides in CSA and the other resides in the DL/I subordinate address space (DLISAS). The PSBUT service can display utilization of either pool.

The CSA pool contains the TP PCBs and the Fast Path PCBs. The DLISAS private-area pool contains the full-function PCBs. This division of pools results in a split of approximately 20 percent in CSA and 80 percent in private storage.

Monitoring Pool Utilization

The following Resource Monitor services can be used to monitor usage thresholds for these pools:

PIPL	Program Isolation
WKAP	General Work Area WKAP pool
DBWP	Database Work Area DBWP pool
PSBW	PSB Pool
DMBP	DMB Pool
MFSP	MFS Pool

If thresholds for these pools are consistently exceeded, consideration should be given to increasing the appropriate pool. Keep in mind that increasing pool sizes may strain another resource such as real storage (paging) or CSA.

Suggestions

- Check the DMBUT display over time to see if the pool is remaining stable. If not, DMBs may be washing out of the pool and incurring large overhead. This can be investigated by using the IMS log tape utility to check for log record types 20 and 21 (database open/close).
- To reduce fragmentation and the possibility of a pool space failure, heavily used DMBs should be made resident. Heavily used PSBs and the intent lists should also be made resident. This causes them to be placed in separate areas, not in the pools.
- Use the DOPT option of the APPLCTN macro only on a test system. It results in a directory search of ACBLIB to locate the PSB every time such a program is scheduled.
- Access to the ACBLIB may affect scheduling and thus system performance. Many of the suggestions made in “Message Format Service (MFS)” on page 4 regarding the MFS format library are applicable.
- The following points should be considered to help reduce program fetch activity:
  - The default number of entries in the BLDL list (dynamic entry count) is 20. This can be increased at region startup with the parameter `DBLDL=nnn` to reduce directory searches. It should be set to zero on a test system to ensure that the newest copy is always fetched.
  - Preloading is especially effective for small, frequently used subroutines (for example, COBOL and PL/I subroutines). It is not necessary for these modules to be link edited as re-entrant or reusable, but they must be so logically. Paging probably would adversely affect large application programs that are preloaded.

– The steplibs for each dependent region should be concatenated with the most frequently used application program libraries first and the IMS RESLIB last. The application program libraries should be full-track blocked to reduce reads.

– Consider using the OS/390 VLF feature to manage program loading.

– Analyzing Operational Displays

There are six standard operational displays of data concerning the functioning of the IMS dependent regions. The *MAINVIEW for IMS Online – Customization Guide* describes how to modify displays or include others tailored to site requirements. In addition, the system status display (STAT) provides an overview of the status of the total IMS.

These comprehensive region displays show the work that the IMS system is currently performing. They should be used to monitor this activity regularly so that problem areas can be located as they occur (for example, a program loop in a dependent region). These displays also show up to eight regions concurrently so that activity in different regions can be compared.

Total System Monitoring

The system status display (STAT/STATR) is designed to simplify continuous monitoring of the system, pointing out possible problem areas that can be investigated in detail with the other displays, or confirming that all is running as smoothly as it should. If there are standard recognizable problems, warning indicators are set and highlighted. Other problem areas may be recognized only in light of knowledge about specific operating characteristics of a particular system. Therefore, the data displayed should be scanned even when no warnings are shown. If the asynchronous services option is available, many variables can be monitored automatically and warning messages sent when user-defined thresholds are exceeded.

There are two basic parts to this display:

- An overview of the critical resources that allow work to be accomplished
- An overview of work already performed and still to be done (that is, transactions processed)

The dependent regions are also considered to be resources of the system (and the most important). If one or more regions are not performing as they should, many resources probably are being misused (that is, pool space, I/O, CPU time, and the like), and throughput will suffer. The region displays are needed to analyze the actual work being done and the load balance. STAT and STATR are meant to be used only to check status, and the STAT/STATR information is very condensed to show 15 regions at a time.

The warning indicators generally suggest the use of another display. For example, if a region is in a PI wait (`WARN=W-PI`), the PI display shows what resource is needed and what region is holding it. However, these are warnings only, not necessarily actual problems. It is natural to have waits, so it is only when such a condition continues (through several refreshes of this display) that further analysis is indicated. If it continues, use DREGN (Program Isolation Activity area of display) or the IRLM display service to see if the condition is valid for the program being processed or a real problem. Workload Analyzer service DWAIT is an effective service to use to determine if this is a persistent problem.

The NOBK indicator is of special value when the system is in trouble and one or more regions should be cancelled. If this indicator is on, a cancel causes IMS to come down also. Previously, there was no way to know if regions were in this state.

Although the dependent regions are highly important, other resources also can be critical. If the PI pool has insufficient space, dependent regions can abend with a U775. By checking for the THRESHOLD warning on the percent ALLOCATED field regularly, steps can be taken to reduce the utilization before it causes a problem.

The line showing stopped resources on the STATR service points out unexpected conditions. If any of these values is higher than expected, the /DISPLAY STATUS command allows further analysis.

The lower part of the STAT/STATR display shows the overall performance of the system:

- How many transactions have been processed since restart?
- How many transactions are currently queued for processing?

These two counts are the true indicators of how well all resources are being used and if they are sufficient to handle the load. They can be analyzed only in comparison to what is expected of a particular system. It is possible to calculate the rate at which transactions are being processed by making two observations, calculating the difference, and dividing by the time difference. If this is done regularly at different times of the day, a very valuable system profile can be built. This profile can be used to spot deviations or slow degradation and to monitor how the system is handling a variable or increasing workload. The rate of transaction arrivals can be calculated by tracking the differences over time in the sum of the current queue count (Q'D) and total transactions processed (PROC). The Resource Monitor monitors of transaction arrivals and transactions processed do this automatically.

See Part 3, “Monitors” for more information about:

- Monitoring paging (PAGE)
- CSA utilization (CSAUT)
- CSA fragmentation (CSAFR)
- SIO activity (SIO)
- EXCP event (SYSIO and DLIO)
- Logical channel busy (CHBSY)

To obtain a broad perspective of the system workload and to determine where a transaction is spending most of its time, use the Workload Analyzer workload wait services (MWAIT/DWAIT). These services identify possible bottlenecks and can provide sufficient information to make a correction. Otherwise, use one of the Resource Analyzer services to analyze the problem further.

Dependent Region Monitoring

The regions display (REGNS) can be used for regular monitoring of the regions. REGNS (described in the “Region Displays” chapter in the *MAINVIEW for IMS Online – Analyzers Reference Manual*) contains information about the current activity in each IMS dependent region. From this display you can see:

- Which application programs are active
- How much work they have done (message queue and database access)

You can then use the DREGN display to see:

- Approximately how much work the application programs still have to do (how many transactions of this type are still queued)
- The specified processing limit

If you recognize a problem, check the logical terminal name in either the REGNS (Summary View) or DREGN display to identify the user who entered the IMS transaction being processed.

If an application program remains in a region for a suspiciously long time, a loop may be indicated. Use the REGNS Message View, DL/I View, and DB2 View to look at message queue and database activity to help locate the problem area. If the number of input messages queued (QUEUED) or dequeued (M-DEQ) is large, perhaps only the processing limit definition (PRLIM) needs to be changed.

Response time to users can be investigated by using REGNS or DREGN to view transaction elapsed time (TRN-ELAP or ELAPSED). An increase in this time is a danger signal indicating performance degradation.

Often information in a region display indicates possible problems that can be checked out more thoroughly with other displays. For example:

- If a program loop is suspected and the total database calls (TOT shown by the DL/I view of REGNS) is high, more specific information on the types of calls being made can be seen in the DLIST or DREGN display.
- If a large number of program isolation waits is shown by the Program Isolation Activity area of the DREGN display, look at the PI service display.
- If regions are often idle, the distribution of input transactions in the class queues can be investigated in CLASQ.
- If scheduling problems are suspected, the number of scheduling failures and their causes can be investigated in the scheduling display (SCHED).

Program isolation is an important automatic feature of IMS to avoid interference between programs in database access and updating by enqueuing on database records. The Program Isolation Activity area of DREGN shows the total enqueues, dequeues, PI waits, and the number of current enqueues for each program currently processing. From this display, it is possible to see if one program is using all the resources (excessive total or current enqueues) and impacting other programs by causing PI waits.

The DL/I View of REGNS or the DL/I Call Activity area of DREGNS can be used by the database administrator (DBA) or anyone concerned with application program performance. For example, these services can be used with a test system to monitor the activity performed by a new application program.

The information shown by the System Activity area of the DREGN display applies to dependent region OS/390 data. The data elements describe the region, not just the current IMS application (for example, in a message processing region where many transactions are processed by various programs). This information identifies whether the corresponding IMS task is currently executing in the dependent, DL/I, or control region (some processing, such as database access, usually occurs in the DL/I region). It shows each address space and its OS/390 dispatching priority and position in the dispatch queue.

The System Activity area of DREGN also shows the current swapping status of the region. IMS dependent regions are marked nonswappable when they are started. A small number of swaps can occur before this is completed. In addition, the total elapsed time each region has been up is shown in hours:minutes:seconds, and the total CPU time and SRB time is shown in seconds. These figures show how well the processing load is balanced between the regions.

The Paging Activity area of DREGN indicates the amount of paging DREGN in each of the dependent regions. The number of page-ins, page-outs, and reclaimed pages are displayed for that time interval, for VIO, and for the common area (CSA and LPA - no pages out). You may want to request this display at regular intervals with the CYCLE SETUP option from the Primary Option Menu to calculate the differences in the amount of this activity at various times of day.

The region transaction profile shown by the PSB/Transaction area of the DREGN display provides additional information about current transactions being processed, including the number currently queued and the total already processed since IMS restart. This information shows possible imbalances in region, class, and transaction assignments. The average length of all input messages received with a transaction code (AVG LENG) is calculated by the IMS queue manager. Tracking these average lengths for high-volume transactions can determine the optimal lengths for the short and long queue data sets (see “Analyzing Queuing” on page 6).

For more detailed analysis of transactions, workload trace (MTRAC/DTRAC) can be used to see the number of calls and database I/Os. A detail trace is useful when testing new application programs to see if they conform to specifications in the number and types of calls.

Analyzing Databases

This section describes OSAM, VSAM, and fast path buffer pools. It also includes information on Hit Ratios and VSAM Hiperspace.

OSAM Buffer Pool

A subpool concept similar to VSAM is employed in which the buffers in any subpool are all the same size. A database is assigned at open time to the subpool with the smallest buffer size in which its blocks fit, or for the specific subpool designated in the DFSVSM*xx* member by the database administrator.

To further improve the speed of searches to locate records in a subpool, the buffers are chained off an array with the same number of entries as buffers. The particular anchor point is determined by hashing of the DMB/DCB/RBN. This substantially reduces the number of buffers that must be examined. Because an increase in pool size does not cause a comparable increase in the number of buffers to be examined, this is particularly effective for large subpools. This allows increase of the pool size as long as the paging rate stays down.

Finally, the method of selecting a buffer to be reassigned (buffer steal) has been altered. During the selection process, each buffer in a subpool is considered to be at 1 of 11 levels (0-10), depending on the work necessary to free it. For example, empty buffers are assigned to level 0 or 1, buffers that are not currently available are at level 9 or 10. Then the subpool is divided into limited search groups with the number of buffers in each group being twice the number of scheduled regions (or one during emergency restart).

A complicated progressive search algorithm is then used to select a buffer, if one is available. The factors considered are the buffer level, the previous owner (the requestor's own buffers are preferred), the least-recently-referenced order of the buffers, and the length of the search. The net result of all this is to reduce the search time even in large pools. The BUFFERS STOLEN and TOTAL SEARCHED by level can indicate the success of this algorithm. Levels 0, 2, and 4 show steals of the requestor's own buffers; levels 1, 3, 5 belong to another.

To make optimal use of these features, you may need to modify the database DBDs and an unload/reload. In particular, the block sizes should be as close to the defined buffer sizes as possible without exceeding it.

Once the block sizes of the IMS databases are known, the number of buffers and buffer size of the subpools can be determined. At data set open time, the first subpool whose buffers are large enough to contain the blocks is assigned. It may be possible to use this feature to separate data set blocks into different subpools. In particular, the index blocks can be isolated from data blocks. However, you must closely investigate the desirability of this.

Once you have chosen the buffer lengths of the subpools, determine the number of buffers. MVIMS can be useful in evaluating database performance. The number of OSAM WRITES-STEAL per second is a good indicator for optimizing the number of buffers in a subpool. If it is very low, there probably are too many buffers, which wastes memory and may increase the paging rate. If it is high, the number of buffers probably should be increased. STEAL WRITES are especially expensive because the buffer handler may issue an IMS log tape WRITE AHEAD call to guarantee that database changes have been physically written to the log before actually writing to the data set. This can result in two additional waits before the buffer can even be reassigned.

VSAM Buffer Pool

The IMS VSAM buffer handler/pool consists of 0 to 255 subpools of fixed length buffers. At database open time, the smallest subpool that can contain the control intervals is assigned, unless the database administrator has assigned a specific subpool for that database.

Although this fixed length buffer subpool technique (similar to current IMS OSAM buffering) eliminates pool compactions, the buffers in each subpool are chained in a least-recently-referenced use chain. This causes all buffers to be examined during searching. When a buffer is needed, VSAM selects the least-recently-referenced. If the buffer contains data that has been modified, it is first necessary to write it out. To improve performance, VSAM writes are deferred until space is needed or an explicit request is made to purge a user's buffers.

The number of WRITES (VSAM initiated) per second should be monitored to see if sufficient buffers have been allocated. Such writes are expensive because they may require a log tape WRITE-AHEAD call. This can result in two additional waits before the buffer can be used.

To prevent this condition, IMS employs a BACKGROUND WRITE feature, which can be turned on or off at system initialization (OPTIONS statement). This is a low priority IMS ITASK that runs when VSAM notices that the next buffer to be freed on the use chain has been modified. It forces VSAM to write out a specified percentage of buffers from the least-recently to the most-recently-referenced. The NUMBER OF TIMES BACKGROUND WRITE INVOKED per minute may also give an indication if there are enough buffers.

Note: Various internal traces may introduce overhead to DL/I calls and buffer handling. The status of these traces may be observed in the Resource Analyzer service and may be turned on or off by the IMS /TRACE command.

The fixed length buffers make it possible to separate control intervals by subpool. Of special importance is the separation of index blocks from data blocks.

IMS supports an index-only VSAM buffer pool. This allows separation into separate pools data and index control intervals of the same size.

IMS supports buffer pools that are dedicated to specific databases. This can be used to give preferential treatment to these databases or to isolate high activity databases.

Hit Ratios

Buffer pool hit ratios can be used to evaluate how well the buffer pools are performing. The hit ratio is the percentage of buffers that were found in the pool without needing to access external storage. Generally it is recommended that data buffer pools have a hit ratio of 60 or higher and index buffer pools have a hit ratio of 80 or higher. These hit ratios are available with the following services:

Table 1. Services Supporting Hit Ratios

	Analyzers	**Monitors**
OSAM	DBST	DBHIT
VSAM	VSST	VHIT

VSAM Hiperspace

When defining your VSAM buffers to IMS, you can request that VSAM allocate Hiperspace buffers in expanded storage to augment your virtual storage buffers. Buffers that normally would have been washed out of the pool due to the least-recently-used algorithm will be migrated instead to Hiperspace buffers. This allows you to have more buffers without increasing your virtual storage requirements.

The number of successful reads from Hiperspace is a measure of the number of reads with I/Os that were saved by using Hiperspace. The number of unsuccessful reads represents the number of Hiperspace buffers that were stolen because of a shortage of expanded storage. To get the full benefit of your Hiperspace buffers, the number of successful writes to Hiperspace should be less than the number of successful reads from Hiperspace. If the number of successful writes is greater than the number of successful reads, buffers are being written but never read.

The Hiperspace hit ratio shows the portion of the VSAM hit ratio contributed by Hiperspace. In other words, your VSAM hit ratio would be smaller by the Hiperspace hit ratio amount if you had not defined Hiperspace buffers. This is displayed on VSST and monitored by HPHIT.

Fast Path Buffer Pool

Buffers from the fast path buffer pool are fixed per dependent region that requires buffers. This is determined by the `NBA=` keyword in the dependent region startup JCL. If the buffers are not available in the CICS environment, the region fails to connect to IMS (either DBCTL or IMS DB/TM). Altered buffers are not written to the area data sets until synchronization point.

The keywords that define the buffer pool in the control region JCL are:

- BSIZ, which specifies the size of the largest DEDB Control Interval (CI)
- DBBF, which defines the total number of buffers
- DBFX, which defines how many of these buffers are to be used as reserve buffers

These reserve buffers act as a cushion at synchronization point time to allow the asynchronous output threads some time to complete writing the buffer while the next transaction is processing. If the output threads are being delayed or there are not enough of them (OTHR keyword), a wait occurs. IMS does not keep global statistics on this wait condition and it must be analyzed by examining the IMS PR report TRNFP or the batch fast path log analysis utility (DBFULTA0).

Note: A buffer is fixed in the buffer pool for each active area that has a sequential dependent part defined (SDEP).

Dependent region allocations of buffers are determined by the NBA and OBA keywords. The normal buffer usage (NBA) should be large enough so that overflow buffers (OBA) are seldom used. IMS reserves one set of overflow buffers out of the buffer pool that is equal to the largest OBA specification in all of the dependent regions. This means that only one dependent region can use these buffers at a time causing serialization; for example, contention for overflow buffers. Detailed analysis of the dependent region buffer usage is performed by the IMS PR report TRNFP or the IMS batch utility (DBFULTA0).

The FPBST display assists in online determination of buffer usage. By refreshing this display and monitoring the buffers in use by region, you can detect an NBA specification that is too small; for example, in the region detail display the USED is consistently more than the NBA.

The OTHR and DBFX values can be tuned by monitoring the number of IDLE OUTPUT THREADS and FIXED BUFFERS AVAILABLE. If the idle threads are consistently 0, they should be increased. If the fixed buffers available are consistently a small percentage of fixed buffers, there is a strong possibility that waits will occur for buffers. To overcome waits due to insufficient output threads or available fixed buffers, increase the number of pre-fixed buffers (DBFX) and/or the number of output threads (OTHR).

Points Common to All Buffer Pools

In evaluating IMS system performance, you should determine the amount of physical I/O per second. Although it is possible to break the activity down by subpool with Resource Analyzer and Resource Monitor services, the PERFORMANCE REPORTER offline Database I/O report can be used to determine activity by DMB. Using the DC Monitor, it is possible to determine activity by data set; however, even this does not show contention between data sets on the same device or activity to a data set spanning multiple volumes (devices). This information can be obtained by using CMF MONITOR or MAINVIEW for OS/390, both BMC Software products, or the IBM Generalized Trace Facility (GTF). These give a better indication of the actual use or contention of devices and paths. This also can be estimated by analyzing the database change records. However, this method is incomplete because retrieve type calls are not logged and a special analysis program would have to be written.

The number of events per second reported, using the techniques in the previous paragraph, do not take into account the number of transactions that generated the activity. The number of events per second could be re-expressed in terms of events per PURGE by using either the VSST global display or DBST display service. A more detailed investigation can be made with the MVIMS transaction processing report.

These calculations should be made at different times (when the system is busy and when it is not busy). Investigate changes observed over time that show a trend. In particular, an increase in READS/SEARCHES per PURGE shown by the VSST global display or DBST display service may indicate the searching of an overflow chain. This could be an HISAM/HIDAM-INDEX overflow chain or an HDAM synonym chain. If so, the database involved should be determined and reorganized.

Suggestions:

- HDAM is probably the most efficient IMS access method; however, long synonym chains must be watched. If encountered, investigate:
  - Randomizer - the most widely used and efficient randomizer is the WORLD TRADE (DFSHDC40).
  - Insufficient RAPs for the number of records (resulting in long synonym chains).
  - Too many RAPs per block for the record size.
  - Not using the BYTES option. This is especially devastating after a reorganization.
  - The advantage of distributed free space.

HISAM is the best for sequencing small records.

  - In VSAM, watch CI/CA splits. It may be possible to use the VSAM KSDS options to reduce these splits.

HIDAM seems to be the most used or abused.

  - The HIDAM INDEX is nothing more than a HISAM database and suffers the same problem of long overflow chains or CI/CA splits.
  - If much sequential processing is done, it should have both forward/backward pointers at the root to reduce index access.
  - If many dependent segments are added to a record, it may help to specify distributed free space.

- IMS data sets should be allocated by cylinder as a single extent. If necessary, they should be reorganized periodically to a single extent. This can be done with IEBGENER for OSAM. EXPORT/IMPORT or REPRO can be used for VSAM (this reorganizes a KSDS also). However, this is not a substitute for a complete IMS database reorganization.
- Watch the placement of data sets to prevent an over- or under-utilization of a device, control unit, or channel. Try to avoid placing two very active data sets on the same device. Also, watch for an excessive number of alternate tracks, especially when these tracks are associated with a high usage data set. This situation tends to reduce seek time.
- Write check should not be specified for IMS databases. There is no integrity exposure because the data sets can be recovered using the IMS utilities and log tapes.
- For VSAM, SPEED should be requested because IMS does not use VSAM recovery. This speeds up initial loads.

Analyzing IMS Internals

This section describes latches, logs, pools, and program isolation analysis.

Latches

IMS latches, like OS/390 locks, are used to protect the integrity of certain resources in a multi-programming environment. A certain number of conflicts are to be expected and show that the latches are performing a needed function. However, excessive conflicts may indicate a malfunction in system performance that you should investigate.

In general, latches are held for only a very short time, which reduces the number of conflicts. If something occurs that increases the time a latch is held, the probability that another task will need the protected service before the task completes also increases. The main causes of such delays are page faulting or faulty dispatch priorities. If a latched routine has to page fault through a pool, the execution of any other task waiting for that service or resource can be blocked, causing a degradation in IMS performance.

Suggestions:

- Investigate the paging rate of the system and consider page-fixing the affected pools (use the LATCH service to see where conflicts are occurring most frequently).
- If latch conflicts are appearing for DMB user routines, check these routines for waits or possible page faults.
- If the conflicts are frequent for certain OSAM buffer subpools, review the definition of the number of subpools and the number of their buffers to reach a better balance between the allocation of the available space and the actual usage of the buffers.

Logs

Three IMS data sets are used, two online and one offline. Resource Analyzer service LOGST displays utilization information and statistics about the two online data sets and provides certain restart data.

Online Log Data Set (OLDS)

The IMS online log data set (OLDS) can be either single or dual and contains complete log records.

IMS only writes the buffer (padded if necessary) when the buffer is completely full. If OLDS and WADS are allocated by JCL, the number of buffers for the OLDS is specified on the OLDS DD statement. If OLDS and WADS are allocated dynamically, the number of buffers is specified by the BUFNO parameter in PROCLIB member DFSVSMxx.

When an OLDS becomes full, it must be archived. Once the archive batch job is complete, the OLDS can be reused. If dual OLDS are in use, a switch is made when either OLDS becomes full.

Buffer size for the OLDS is taken from the preallocated data set. At least four or five buffers should be assigned because one of the buffers is used to read the OLDS if a dynamic backout is required.

Write-ahead Data Set (WADS)

The IMS write-ahead data set (WADS) is preformatted and is used to write incomplete OLDS buffers. WADS does not have its own buffers; it uses the OLDS buffers. The OLDS buffers are segmented in 2K segments. Any check write requests cause unwritten segments for the current buffer to be written to the WADS. This allows the log write-ahead requirement to be satisfied.

Database log write-ahead is no longer an option; it is compulsory.

Data communications log write-ahead (DCLWA) is a default option. Review it carefully. Response time can be adversely impacted if the default `DCLWA=YES` is allowed or specified. Response time is impacted while the transaction or message is written to the WADS, which delays processing; however, this does give complete DC integrity and should be evaluated.

IMS restart processing accesses the WADS to close the OLDS in the event of a failure.

System Log Data Set (SLDS)

The system log data set (SLDS) is an offline data set and can be either tape or DASD. The SLDS is used for archiving the OLDS.

Log Performance Suggestions

- Prevent DASD contention on all online log data sets; that is, separate paths, strings, devices (no shared DASD), and so on. Consider contention that can occur because of an archive utility that is executing as a batch job.
- Allocate the WADS on a low-usage device because this is most critical.
- Provide for multiple WADS backups.
- Allocate OLDS and WADS with contiguous space.
- Specify approximately 10 buffers in the archive job for both OLDS and SLDS.
- Ensure that the archive job runs at a high priority.

Note: The Resource Analyzer service LOGST produces a warning message if either the last OLDS is in use or the system is waiting for an archive to be performed.

The LOGST service also displays the OLDEST LCRE. This field identifies the oldest recovery point that IMS will require in the event of a restart. A very old time in this field could indicate that a BMP or JBP is not taking sufficient checkpoints and the record required for restart might already be archived.

Program Isolation Analysis

Program isolation is the IMS function that allows many application programs to access the same databases concurrently without interference. It achieves this by enqueuing on each database record as it is accessed at READ, UPDATE, or EXCLUSIVE level. Normally, each enqueue requires two QCBs (queue control blocks) from the PI pool. To improve performance, an entire record is locked by enqueuing the root segment at UPDATE level when a segment is accessed. When moving from record to record, the new root segment is enqueued at UPDATE level and the old root segment is dequeued. Dependent segments are enqueued only when modified. These enqueues remain until the modifying program reaches a synchronization point. Any other program attempting to access a segment or record that has been so enqueued must wait.

The PI service allows investigation of program isolation problems. Watch the amount of FREE SPACE in the DYNAMIC POOL closely. If more space is required and the MAXIMUM POOL SIZE has not been reached, a GETMAIN is done in subpool 241 (CSA) for the amount specified in INCREMENT. When the pool can no longer be expanded, the requesting program is pseudo-abended with a u775. All changes are backed out and the transaction is put back on the queue for reprocessing. This is an enormous drain on IMS and causes a sharp decline in performance.

However, it is not a good idea to allow the PI pool to grow uncontrollably. Once space is obtained, it is never released. Because of the space management algorithm, the total available space is constantly referenced even after most of it is free. This can increase the page faulting rate for all IMS regions. Because this is also one of the few pools that cannot be page-fixed, page faulting can be reduced only by keeping this pool small. Use the PIMAX monitor to make sure that one region is not using this pool too much. If it is caught early enough, a region can be stopped before it fills up the pool.

The bottom portion of the PI display gives the number of enqueues each dependent region currently holds. This is broken down by level of enqueue:

- READ
- UPDATE
- EXCLUSIVE

More importantly, it is possible to see if a program is in a PI wait, what resource it is waiting for, and which program is holding the resource. If there are frequent PI waits, serious attention should be given to the application system design or database design. This information is especially valuable in determining the checkpoint call frequency needed in BMPs and JBPs. You can possibly reduce contention by changing the processing option to GO or EXCL to bypass PI. Some cases may require multiple database PCBs with different PROCOPTs. You can use explicit program enqueue/dequeue control using Q command codes to control simultaneous access. In a few cases, the problem may not be solved without a complete system redesign.

In addition to lengthening program execution time, PI waits may lead to deadlocks. Although not fatal, IMS must pseudo-abend one of the programs. All of this transaction's changes are dynamically backed out and the transaction is put back on the queue for reprocessing.

Note: IRLM can be used to control database contention in place of program isolation.

Pools

There are two kinds of pools: CBT and non-CBT. The POOLC service displays CBT pools. The POOLS service displays the status of the non-CBT pools. The DPOOL service displays detailed information about non-CBT pools.

CBT Pools

Although you do not have direct control over most CBT pools, you can control DPST and SAP. You should consider the following:

DPST` Pool

The DPST pool holds the dependent region partition specification tables (PSTs). DPST is defined by the Stage 1 system generation, (`IMSCTRL MAXREGN=` or `DFSPRRG0 PST=`); this definition can be overridden by the JCL parameter `PST=`. DPST defines the minimum number of PSTs that IMS should hold available. Always define the average number of regions that you intend to run. Understating the number does not cause a PST shortage because the pool automatically expands and contract. However, the GETMAIN and FREEMAIN processes are an unnecessary overhead and can be avoided by making the correct size definitions.

Additionally, `PST=` defines the number of VSAM strings, which do not change. If there is a shortage of VSAM strings due to an understatement of `PST=`, the dependent region waits.

SAP Pool

The SAP pool is used to hold all the dynamic and other save area prefixes (SAPs). SAP is defined in the Stage 1 system generation macro (`IMSCTRL MAXIO=` or `DFSPRRG0 SAV=`; this definition can be overridden by the JCL parameter `SAV=`. This pool does not expand and its shortage causes IMS to enter selective dispatching which severely degrades performance.

Non-CBT Pools

WKAP Pool

Main pools contain the DFSISMN0 control blocks and a general work pool, WKAP. MVIMS displays only the WKAP portion. IMS creates temporary work pools from WKAP. The size of a temporary work pool is defined by `DFSPRRG0 WKAP=`; this definition can be overridden by IMS control region JCL parameter WKAP.

QBUF Pool

The QBUF pool is the central storage area used as I/O buffers for the three types of message queue data sets: SHMSG, LGMSG, and QBLKS. The IMS storage manager module DFSISMN0 allocates the storage for the QBUF pool during IMS initialization. The IMS message queue buffer manager module DFSQBFM0 manages the QBUF pool during transaction processing. A well-tuned QBUF is essential to fast transaction response time because all inbound and outbound message traffic must pass through QBUF.

QBUF Allocation:

The size of the QBUF pool is defined during Stage 1 system generation by the MSGQUEUE macro. You can use the QBUF parameter in the DFSPB*xxx* parmlib member or in the IMS JCL EXEC statement to override the number of QBUF buffers at execution time. You can use the QBUFSZ parameter in the EXEC statement to override the size of the buffers. To page-fix the QBUF pool, set EXVR=Y in the EXEC statement for the control region JCL.

QBUF Internals and Associated MVIMS Displays:

IMS caches as many messages in the QBUF pool buffers as possible before it attempts a buffer steal, a process that requires I/O to one of the message queue data sets SHMSG, LGMSG, or QBLKS. A buffer steal is not attempted until all the buffers have been used. DFSQBFM0 attempts buffer steals from the stealable buffers first. A buffer is stealable if it is not currently involved in I/O and it has no IWAITs. To steal a buffer for use by a different message, DFSQBFM0 must first write the existing buffer contents to one of the message queue data sets. If a buffer is currently involved in an I/O or otherwise owned by another process (busy, in other words), an IWAIT is always required before the buffer can be written and stolen. For that reason, access time to a stealable buffer is much faster than access time to a busy buffer.

When a buffer is being stolen, the QBSL latch is held. MVIMS displays IWAIT data for the QBSL latch in the LATCH Summary display. The latch may be held without causing an IWAIT. However, if there are no stealable buffers when an attempt to acquire the QBLS is made, an IWAIT for QPWTBFR (wait for an available buffer) is incurred before the latch can be acquired. The number of IWAITS for QPWTBFR is provided in the QUEST Statistics analyzer display in the WAITS FOR AN AVAILABLE BUFFER field. If the value in the field is not zero, all QBUF buffers are busy and none are stealable.

Area 1 of the DPOOL analyzer display for the QBUF pool shows the number of stealable buffers (STEAL) and busy buffers (CURR). The POOLS analyzer display shows the number of busy buffers (CURR), and the value is the same as the CURR value in the DPOOL display. The POOLA monitor monitors the QBUF pool for high contention on its buffers. Contention is defined as buffers that are busy (the CURR value in DPOOL and POOLS) and not stealable (the STEAL value in DPOOL). Using this information, you can see what portion of the QBUF pool is busy (CURR) and what portion is used but not busy (STEAL).

QBUF Tuning:

If your system is not real-storage constrained, your QBUF pool should be set large enough to ensure minimal I/O to the message queues. In general, the RECLENG parameter in the MSGQUEUE macro should be set so that I/O activity to SHMSG and LGMSG is evenly split. The split is determined by message segment lengths and the distribution of message arrival between message segments shorter than LGMSG LRECL and longer than SHMSG LRECL. I/O to the message queues may occur because the mix of arriving message segments uses up one or the other section of the QBUF pool buffer even though the size of the QBUF pool seems large enough. You should carefully balance the path length of buffer searching and I/O to the message queue data sets. If you overallocate the number of buffers, the buffer search path length and CPU cycles will increase. If you underallocate the buffers, I/O will increase.
One tuning technique to reduce storage requirements and CPU cycles for QBUF buffer operations is to decrease the number of buffers until the QUEST Statistics analyzer display shows increasing IWAIT instances for message queue I/O.

PSBW, PSB, and DMB Pools

These pools are managed by DFSISMN0. If `LSO=S is specified`, the PSB pool is split into two parts. The POOLS service shows the parts as DLMP/PSBC in CSA and DPSB/PSBD in private for the DLISAS address space. The DMB pool is shown as DLDP/DMBP in CSA. You can use the DPOOL service to get details for each pool. Each pool has associated monitors that can be started (such as PSBP and DMBP).

During transaction scheduling, IMS allocates space from these three pools to accommodate the required PSB and DMB. If space allocation fails in any of these pools, then IMS tries to make space available by purging not-in-use control blocks. If space is still not available, then scheduling fails and IMS tries to schedule the next eligible transaction. The dependent region waits if none can be scheduled. For more information, see "Pool Utilization" on page 10.

DMBW is the DMB work pool. Although it is not used during transaction scheduling, it is used later by the DL/I delete/replace action module (DFSDCDC0).

The resident PSB and DMB are loaded at IMS startup time into storage outside these pools. The resident PSB is copied into the PSB pool when needed. The resident DMB does not need to be copied.

The PSB pool holds the PSBs on a most-recently-referenced basis. It is defined in the Stage 1 system generation (`BUFPOOLS PSB=`, `SASAPSB=` or `DFSPRRG0 PSB=`, `CSAPSB=`, `DLIPSB=`); this can be overridden by the JCL parameters `PSB=`, `CSAPSB=` and `DLIPSB=`. The IMS default for splitting the PSB pool (80/20) is reasonable, so let it default.

The PSB work pool (PSBW) holds various DL/I work areas for the PCBs such as index, SSA, and SPA. It is defined in the Stage 1 system generation (`BUFPOOLS PSBW=` or `DFSPRRG0 PSBW=`); this definition can be overridden by the JCL parameter `PSBW=`.

The DMB pool holds the DMBs on a most-recently-referenced basis. Space shortage causes some DMBs to be closed and flushed. The DMB pool is defined in the Stage 1 system generation (`BUFPOOLS DMB=` or `DFSPRRG0 DMB=`); this definition can be overridden by the JCL parameter `DMB=`. The DMB work pool is defined in `DFSPRRG0 DBWP=`; this definition can be overridden by the JCL parameter `DBWP=`.

CESS Pool

The CESS pool allocates external subsystem communication control blocks, such as the DB2 attach facility.

MFBP Pool

The MFS buffer pool, MFBP, accommodates MFS control blocks. However, MFSTEST control blocks are not taken from this pool; they are taken from the CIOP pool. MFBP size is defined in the Stage 1 system generation, (BUFPOOLS `FORMAT=` and `FRE=` or in `DFSPRRG0 FBP=` and `FRE=`); this definition can be overridden by the JCL parameters `FBP=` and `FRE=`. This pool is allocated from extended CSA.

CIOP, HIOP and RECA Pools

The communications I/O pool, CIOP, is used as message buffers between VTAM/BTAM and the IMS queue manager. RECANY has its own CBT pool, RECA. CIOP contains only the output buffers and EPCB blocks, described below.

A high communications I/O pool, HIOP, is allocated from IMS control region extended private area to expedite the CIOP usage. This pool is defined in the Stage 1 system generation (macro `BUFPOOLS COMM=` and macro `COMM RECANY=` or in `DFSPRRG0 TPDP=`); this can be overridden by the JCL parameter `TPDP=`.

Because the MFSTEST work area is taken from the CIOP pool, consider the sizes of both MFSTEST and the CIOP pool. MFSTEST size is specified in IMS JCL `MFS=`.

EPCB Pool

The EPCB pool holds the Fast Path PCBs. EPCB storage shortage causes the Fast Path transaction scheduling to fail without any indication to you.

Analyzing the System

This section describes dispatching and real storage.

Dispatching

The dispatcher statistics display (DSPST) gives an overview of the status and activity of both OS/390 and IMS dispatching.

In the first section, the IMS control region, the DL/I SAS, and the dependent regions are listed with pertinent OS/390 data. The default dispatch priority is set in the supplied PROCs at 239. The dependent regions should run at a slightly lower priority. To assist IMS in its attempt to balance the activity, the priorities of the dependent regions normally should be equal or quite close. This should be changed only if specific processing characteristics of the installation require special consideration.

The SRM parameters of domain and performance group/period and the current swap status can give an indication of how OS/390 functions are affecting the performance of the IMS regions. As a general rule, only one performance period should be defined for the IMS performance group(s), because the IMS regions should be allowed to do their job without excessive interference from the OS/390 SRM facilities.

The IMS dispatcher is responsible for the IMS internal multi-tasking by creating and dispatching ITASKS. Each ITASK is associated with an ECB (event control block) that is posted when work is to be performed for a particular function and with a SAP (save area prefix) that controls the IMS resources associated with the task (for example, a save area set for the registers of all the invoked IMS routines).

There are pre-assigned SAPs for activity associated with each dependent region, logging, and the like. Dynamic SAPs are used for all terminal I/O activity in IMS/DC. The number available depends on the specification in the MAXIO statement during the generation or the SAV parameter at execution. Because each SAP is associated with a GETMAINed area in CSA of over 1200 bytes for the save area set, this number should be optimized to conserve space.

This can be accomplished by reducing the number of dynamic SAPs each time IMS is brought up (but only until a few occurrences of TOTAL ITASKS WAITED FOR A SAP appear). After this point, performance might be affected by further reduction.

ITASKS are created and dispatched for all terminal I/O activity, and destroyed upon completion. By subtracting the number of ITASKS created from those dispatched, an approximation of the number of IWAITS can be found. This is an indication of the total IMS I/O activity. For example, a dependent region must IWAIT for database access. This number also can be calculated per second and used as a standard performance indicator to be checked over time.

Real Storage

Probably the greatest single cause of poor IMS performance is an overcommitment of real storage. This results in a high demand paging/swapping rate. Not only does it waste I/O, but also CPU. These effects are particularly devastating to a lightly loaded IMS system.

On such a system, the IMS modules constantly page in and out. Because of the design of IMS and OS/390, a region is effectively dead during page fault processing. Once it regains control, it usually immediately page faults again because of the very large working set size of IMS.

The IMS latches and commonly referenced database records are of particular concern.
The only thing that can be done to reduce such interference is to reduce the page faulting in IMS. While it helps to reduce the non-IMS jobs that run concurrently, there may not be any alternative to page-fixing a portion of IMS.

The RS service allows the investigation of real storage usage. Because IMS is such a heavy user of CSA and LPA, these are also included. It gives a complete breakdown of the status of the usable page frames in the system:

- Total
- Pageable
- Fixed
- Long-term fixed

In an identical form, the status of the page frames used by IMS are given.

Chapter 2. User Techniques for Monitors

This section provides assistance in using the functions available with the data collection monitors. This section is not a step-by-step description of all the uses of these functions, because installations have varied requirements depending on their processing load and available resources (which may vary over time) and because there are so many services available.

The built-in flexibility of the monitor services allows you to determine which items are important to observe, when to observe them, and how long to observe them. This flexibility also allows you to vary these specifications easily when needed.

The monitor services offer continuous monitoring, early warnings, graphics, and rate calculations. They can be used as:

- An operations monitor of current IMS events
- A performance analysis tool to assist in tuning the system to use the available resources better and to plan for the future

With the monitor services and the timer facility that drives them, these tasks can be done automatically and selectively.

The following sections describe uses of the monitor services by different people within the IMS organization. This is not a complete list and many of the services could be of value to other people within the organization. The technique descriptions use sample block requests delivered with the product. The descriptions are grouped by the following users to present the various ways that an organization can use MVIMS:

- Master terminal operator
- IMS manager
- IMS performance analyst and system programmer
- Database administrator

Master Terminal Operator (MTO)

The master terminal operator is responsible for controlling the system and keeping it up and running smoothly. Responsibilities include controlling the resources and solving operating problems. The sooner problems are detected and the causes analyzed, resolving or minimizing the effect, the better service IMS can give to the user. The automatic warning feature of MVIMS can relieve the MTO of the chore of continuously monitoring all the different aspects of the system. A set of standard monitor requests can be defined for the installation that can automatically monitor the system and inform the MTO when a potential problem (as defined for that system) is detected.

Warnings are sent automatically to the BBI-SS PAS Journal log, where they can be browsed. Warnings can also be sent to the MTO console. If MAINVIEW AutoOPERATOR for IMS is installed, warning messages can automatically invoke AO EXECs to take action or put the warning message in the AO STATUS/EXCEPTION panel.

Two types of measurements can be set up; one to monitor specific problem areas (such as the usage of a limited resource) and one to monitor general system performance indicators (such as input queue length or scheduling failures).

BLKMTO Example

```
REQ=INQTR  WMAX=20  WLIM=5  LOG=ATWARN  WMSG=MTO
```

This sample request monitors the total input queue length every minute, automatically issues a warning message to the MTO when the queue length exceeds the specified critical threshold (`WMAX=20`), suspends warning messages after five are issued (`WLIM=5`), and logs a plot of the accumulated history data whenever a warning message is issued (`LOG=ATWARN`).

A long input queue is one of the first indicators that system problems exist. Something may be interfering with normal processing, causing a backlog of transactions and a degradation in response time. The critical queue length varies for each installation and the threshold that defines the warning condition (WMAX) should be adjusted accordingly. If an installation stops transactions in the queue often (perhaps priority 0 transactions waiting for BMPs), the MTO might choose to monitor only the transactions that can be scheduled (`REQ=IQSCL`). In a Fast Path environment where the expedited message handler is used, the monitor INQBG would be used. Other factors may influence the definition of a warning condition; for example, the submission of batched transactions, which causes a quick jump in input queue length. Warning messages for this expected condition can be avoided by setting the WIF option; for example, `WIF=3` specifies that a queue length greater than the threshold must exist for three intervals before the warning message is issued. When the MTO is informed of a long input queue, other MVIMS displays, CLASQ, the history PLOT of INQTR, REGNS, DREGN, or SCHED can be used to detect the cause of the problem.

```
REQ=CSAUT  WMAX=80  WMSG=WTO  I=00:05:00
```

This request monitors CSA for a usage percentage. CSA is a critical resource for IMS systems with high program isolation activity. This percentage is checked every 5 minutes (`I=00:05:00`), and a message is sent to the system console if it exceeds 80 percent (`WMAX=80`).

```
REQ=LGMSG  WMAX=70  WMSG=MTO
```

This sample request monitors the usage of the long message queue data set and informs the MTO if it exceeds 70 percent (`WMAX=70`). This early warning should give the MTO time to determine the cause and take action before a full queue data set causes an IMS ABEND. For example, if an application program is in a loop and is writing invalid messages to the queue, the MTO could cancel the program.

```
REQ=@RSPC,1  WMAX=3.5  WIN=2  WMSG=MTO
```

This Workload Monitor sample request monitors MPP transaction response time for transactions in processing class 1. For the purposes of this example, it is assumed that class 1 comprises short-running transactions which must have a 3.5-second or shorter response time to meet service level objectives. When the average response time in an interval for class 1 transactions exceeds 3.5 seconds, a warning message is issued. To give the MTO time to react, subsequent messages are issued every two intervals (`WIN=2`) while the condition persists.

Transaction response time may increase for many reasons. For example, a sudden increase in the arrival rate of class 1 transactions can exceed the capacity of available message processing regions to process them, or an increase in the overall system real storage demand may cause unacceptable amounts of paging. The former case can be diagnosed by using the STAT/STATR display or the history PLOT of ARVCL. The latter case can be diagnosed by using the DREGN display (Paging Activity area) or the history PLOT of PAGE.

There are many other resources and indicators which can be measured, from system data set usage, internal pool usage, and paging, to the level of PI enqueues. Depending on installation configuration and activity, some or all of these services can be defined to automatically monitor the most critical areas. Because the number of requests and the sampling interval for each can be set and modified as needed, the amount of MVIMS activity (and resulting system overhead) can be controlled and channeled to fulfill real needs.

IMS Manager

The IMS manager can use monitor services for high-level monitoring of the total system and for selectively tracking special areas.

BLKMGR Example

```
REQ=PRCTR     I=00:06:00     LOG=ATPD     WMSG=MTO
```

This sample request monitors at 6-minute intervals the number of transactions processed and automatically sends an entry to the IMS log at the end of each period (10 intervals = one hour). This up-to-date, graphically presented data, with the processing rate-per-second already calculated for short- and long-term time periods, is available for online access at any time while the request is active. The sampling interval can be set to show a different spectrum, from seconds to hours. If this request is started before transaction processing begins, the total count field shows the total number of transactions processed that day.

```
REQ=ARVPR,ACC+     LOG=ATPD     RANGES=0,10,60,300     WMSG=MTO
```

This sample request monitors the pattern of transaction arrivals for a group of programs (ACC is a qualified name), automatically sends a plot display record to the IMS log at the end of each period (10 intervals), and displays the distribution of the number of arrivals per interval in the ranges 0 to 0, 1 to 10, 11 to 60, 61 to 300. In this case, the arrivals are monitored for a program group; however, arrivals may be monitored for a transaction, transaction group (qualified name), processing class, program group (qualified name), or application program name. The form of the request depends on installation naming conventions. The IMS manager can use this information to discuss system usage and performance and have statistics for the period in question, not just totals or averages for a full day.

The frequency distribution produced indicates how often a certain arrival rate was measured: no arrivals in one minute, 10 per minute, 60 per minute, 300 per minute or over. The highest count ever measured is shown.

```
REQ=@RSPT,AR+     LOG=ATPD     RANGES=2.5,5,7.5,10     WMSG=MTO
```

This Workload Monitor sample request monitors the transaction response time for a group of transactions (AR is a qualified name), automatically sends a plot display record to the IMS log at the end of each period (10 intervals), and displays the distribution of transaction response time in the ranges of 0 to 2.5 seconds, 2.51 to 5 seconds, 5.01 to 7.50 seconds, and 7.51 to 10 seconds. In this case, transaction response times are being monitored for a transaction group; however, response times also may be monitored by processing class, LTERM, region, program, or user. Qualified names are allowed for every option except class. The form of the request depends on installation naming conventions. The IMS manager can use this information to discuss system response time and service level objectives with users and can generate statistics for any period in question, not just totals or averages for an entire day.

IMS Performance Analyst and System Programmer

The person in charge of IMS tuning has a very complex task: determining the best use of available resources to maximize user service. The large number of variables to be monitored, the many parameters that can be adjusted, the interaction of the various internal IMS functions such as queueing and scheduling, the effect of operating system constraints, the continual variations in processing load and applications mix, and the number and size of the resources to be controlled all increase the difficulty of understanding and tuning the IMS system. Offline reports of summarized data often do not give the precise information needed to analyze the causes of current problems or to detect potential bottlenecks in resource usage.

With the monitor services, the performance analyst can selectively monitor only the information currently needed. The amount of output produced and the cost of producing it can be controlled.

Many different resources and performance indicators can be measured as needed, such as pool usage, I/O rates, program isolation activity, input and output queues, paging, and CSA. Many of these can be measured either globally or selectively which allows the collection of exact information for specific problem areas.

BLKPERF Example

```
REQ=DBSTL,2     I=00:10:00     WMAX=100    LOG=ATWARN,ATSTOP

START=10:30:00  STOP=11:40:00
```

This sample request monitors the number of buffer steal writes performed between 10:30 and 11:40 a.m. in the OSAM database buffer subpool 2. The graphic history can be inspected online at any time after 10:30 until the request is specifically purged. In addition, the requestor is informed at the terminal whenever the number of steal writes exceeds 100 in any 10-minute interval. The history plot is written automatically to the BBI-SS PAS Image log for later offline analysis whenever the warning threshold is exceeded and also at 11:40 when the measurements for this request are discontinued.

Several different I/O counters can be monitored, either by IMS function, by using the MFSIO service or the DBSTL service, or on a system level, such as start I/Os by unit or paging rates.

```
REQ=PSBP     I=00:00:30     LOG=ATPD
```

This sample request creates a detailed graphic history of PSB pool usage. The percent allocated is measured every 30 seconds and a plot record written to the log every 5 minutes (30 seconds x 10 intervals = one period).

This is just one example of how the performance analyst can use MVIMS to selectively monitor special problem areas whenever needed. Other, more generalized requests can be set up to always be active, but only write BBI-SS PAS Image log records when warning conditions occur. Analysis of these records shows areas that need more detailed study.

The WARNINGS WRITTEN field on the Timer Facility Statistics display (Primary Menu Option 5) shows whether any warning conditions were detected.

```
REQ=SIO,158   RANGES=10,60,300,1500
```

This sample request tracks the number of successful start I/Os issued to unit 158, which is used by IMS (for example, a database). In addition to the graphic display of each minute's activity and the calculation of rates-per-second, a frequency distribution is accumulated. The range limits defined are for the activity count within an interval, in this case one minute. This corresponds to a distribution of rates-per-second of 1 every 6 seconds, 1 per second, 5 per second, and 25 per second. The last range limit shows the highest value measured. If IMS is run with `LSO=S`, DLIO can be used for data sets allocated in DLISAS.

```
REQ=CLASQ        I=00:05:00   LOG=ATINTVL

START=09:00:00   STOPCNT=50
```

This sample request writes an IMS log record of the Class Queuing display every 5 minutes until 50 are created. This automatic logging of an informational display can create an audit trail of detailed data at a regular interval within a certain time period. These records can be printed with IMRPRINT and analyzed at any time.

Database Administrator

The database administrator often is responsible for monitoring and controlling the performance of the application programs. With several of the monitor services, an automatic process can be activated, which checks for adherence to installation standards and good programming practices. The effects of new applications on the total system and the performance and usage patterns of the application itself can be monitored.

BLKDBA2 Example

```
REQ=DBTOT(MSGRGN01)   WMAX=20   I=00:00:01
```

This sample request for service DBTOT checks the level of DL/I activity occurring in a region. For example, if the first region (always started as job MSGRGN01) is only supposed to process fast transactions issuing a small number of DL/I calls, this request checks every second and issues a warning message to the BBI-SS PAS Journal log if the program being executed issues more than 20 database calls per scheduling.

```
REQ=PIENQ BMPRGNXX   WMAX=250   I=00:00:10
```

This sample request checks every 10 seconds and automatically writes a warning message to the BBI-SS PAS Journal log if the BMP in this region (when it is active) ever has more than 250 outstanding PI enqueues. The DBA can look at the PI display whenever a warning is received to see if this level of PI activity is causing conflicts with other regions at that time. The plot of this request gives a full history of the enqueue levels reached between CKPT calls. This information allows fine tuning of the frequency of CKPT calls needed in that BMP.

```
REQ=PIMAX   WMAX=1500   I=00:00:15
```

This sample request checks every 15 seconds and sends a warning message to the BBI-SS PAS Journal log if any region exceeds 1500 PI enqueues. The region name and PST number are included in the warning message. This can be used to check if a program is taking checkpoints frequently enough.

Part 2. Using MVIMS

This section describes how to use MVIMS. It includes:

- A description of analyzers, monitors, and traces and how they are used
- The online functions that can be used with MVIMS in a terminal session (TS)
- A description of the Primary Option Menu that provides easy access to product service applications

Chapter 3. The MAINVIEW Product Family

MAINVIEW is an integrated family of performance management and automation products that monitor and control the multivendor enterprise information system. MAINVIEW consists of performance monitors, network integration software, automated operations, and prewritten automation applications.

MAINVIEW product integration allows host and network system monitoring and automation (even in remote locations) through a common user interface, the MAINVIEW Selection Menu, which is shown in the *Using MAINVIEW* manual.

The integration of MAINVIEW products is provided through BMC Software Intercommunications (BBI) technology. For more information about the BBI architecture, see the *MAINVIEW Common Customization Guide*.

Chapter 4. Measuring IMS Activities and Resources

The MVIMS analyzer and monitor services measure these IMS activities and resources:

- Terminal I/O
- IMS queuing
- Scheduling of application programs (PSB and DMB pools) in the dependent region
- Application program activity in the dependent regions
- Database activity and buffer pool utilization
- IMS internal functions and interactions with OS/390
- IRLM functions
- IMS workload

Table 2 on page 40 groups the analyzer and monitor services by the IMS activity or resource area measured and indicates where the services are described. The area indicated in the table is shown as it appears in the AREA field of the service list applications.

For techniques about how to use the services to tune your system, see Part 1, “Performance Analysis and Monitoring Techniques” on page 1.

Table 2. IMS Activities and Resources Measured

IMS Activity/ Resource Area	**Analyzer Displays (Analyzers Reference Manual)**	**Monitors and Traces (all references in this manual)**
Terminal I/O (MFS area)	"MFSST - MFS Statistics" "MFSUT - MFS Pool Utilization"	"MFSFD - Percentage of MFS Blocks Found in Pool" on page 119 "MFSIO - MFS I/O" on page 119 "MFSIR - MFS Immediate Requests" on page 120
IMS queuing (QUEUE area)	"QUEST - Queue Statistics"	"INQBG - Input Queue Length by Balancing Group (BALG)" on page 122 "INQCL - Input Queue Length by Class" on page 122 "INQTR - Input Queue Length by Transaction Code" on page 123 "IQSCL - Schedulable Input Queue by Class" on page 123 "LGMSG - Long Message Queue Percentage Utilization" on page 124 "QBLKS - Queue Blocks Percentage of Utilization" on page 124 "QIO - Queue I/O" on page 124 "QWAIT - Queuing Waits" on page 125 "SHMSG - Short Message Queue Percentage Utilization" on page 125 "INLK - Input Messages by Link" on page 126 "OQLK - Output Queue by Link" on page 126 "OUTLK - Output Messages by Link" on page 127 "DEADQ - Dead Letter Queue Count" on page 128 "OQLN - Output Queue Length by Line" on page 129 "OQLT - Output Queue Length by LTERM" on page 129 "OQND - Output Queue Length by Node" on page 130 "OUTLN - Messages Output by Line" on page 130 "OUTLT - Messages Output by LTERM" on page 131 "OUTND - Messages Output by Node" on page 131

Table 2. IMS Activities and Resources Measured (continued)

IMS Activity/ Resource Area	Analyzer Displays (Analyzers Reference Manual)	Monitors and Traces (all references in this manual)
Scheduling of application programs in the dependent region (SCHED area)	"BALGQ - BALG Queuing" "CLASQ - Class Queuing" "PSBUT, DMBUT - PSB and DMB Pool Utilization" "SCHED - Scheduling Statistics"	"ARVBG - Transactions Arrivals by Balancing Group (BALG)" on page 133 "ARVCL - Transaction Arrivals by Class" on page 134 "ARVPR - Transaction Arrivals by Program" on page 134 "ARVTR - Transaction Arrivals by Transaction Code" on page 135 "PRCBG - Transactions Processed by Balancing Group (BALG)" on page 135 "PRCCL - Transactions Processed By Class" on page 136 "PRCPR - Transactions Processed by Program" on page 136 "PRCTR - Transactions Processed by Transaction Code" on page 137 "SCHFL - Scheduling Failures by Type" on page 138
Application program activity in the dependent regions (REGN area)	"DLIST - DL/I Call Status" "DREGN - Region Detail (Event Collector Data)" "DREGN - Region Detail (No Event Collector Data)" "REGNS - IMS Regions (Event Collector Data)" "REGNS - IMS Regions (No Event Collector Data)" "STAT/STATR - System Status"	"D2CON - IMS Region Connection to DB2 Subsystem" on page 139 "D2SON - DB2 Sign On by Subsystem" on page 139 "D2THD - Active IMS Region Threads to DB2 Subsystem" on page 140 "DBGU - Database Calls per Message Get Unique by Region" on page 140 "DBTOT - Database Calls per Scheduling by Region" on page 140 "MSGGU - Message Calls per Message Get Unique by Region" on page 143 "MSGT - Message Calls per Scheduling by Region" on page 143 "WAIT - Region in a Long PI Wait" on page 144

Table 2. IMS Activities and Resources Measured (continued)

IMS Activity/ Resource Area	**Analyzer Displays (Analyzers Reference Manual)**	**Monitors and Traces (all references in this manual)**
Database activity and buffer pool utilization (DB area)	"DBST - OSAM Global Pool Statistics" "DBST - OSAM Subpool Statistics" "FPBST - Fast Path Buffer Pool Statistics" "VSST - VSAM Global Pool Statistics" "VSST - VSAM Subpool Statistics"	"DBIO - Database I/O Count by Subpool" on page 145 "DBHIT - Hit Ratio for OSAM Buffer Pool" on page 146 "DBSTL - Database Buffer Steals by Subpool" on page 147 "HPACC - Hiperspace Access by Subpool" on page 147 "HPHIT - Hiperspace Hit Ratio by Subpool" on page 148 "HPSTL - Hiperspace Buffer Steals by Subpool" on page 149 "SBUSE - Sequential Buffering Storage by Region" on page 150 "VDBIO - VSAM Database I/O by Subpool" on page 150 "VDBWR - VSAM Writes by Subpool" on page 151 "VHIT - VSAM Hit Ratio by Subpool" on page 152

Table 2. IMS Activities and Resources Measured (continued)

IMS Activity/ Resource Area	Analyzer Displays (Analyzers Reference Manual)	Monitors and Traces (all references in this manual)
IMS internal functions (INTNL area)	"APPCA - APPC Activity Summary" "APPCL - APPC LU Status" "DAPPC - Input Allocation Direction" "DAPPC - Output Allocation Direction" "DLTCH - Latch Detail" "DPOOL - Detail Pool (Non-CBT Variable Pool)" "DPOOL - Detail Pool (Non-CBT Fixed Pool)" "LATCH - Latch Summary" "LOGST - Log Statistics" "PI - Program Isolation" "POOLC - Pool Summary (CBT)" "POOLS - Pool Summary (Non-CBT Variable and Fixed Pools)"	"DBWP - Database Work Area Pool Percentage of Utilization" on page 153 "DMBP - DMB Pool Percentage of Utilization" on page 153 "DSAP - Dynamic SAP Percentage of Utilization" on page 154 "EPCB - EPCB Pool Percentage of Utilization" on page 154 "LAWT - Average Latch Wait Time" on page 155 "LMAWT - Maximum Average Latch Wait Time" on page 156 "MFSP - MFS Pool Percentage of Utilization" on page 156 "OBUFW - OLDS Buffer Waits" on page 157 "OCHKW - OLDS Check Writes" on page 157 "PIENQ - Program Isolation Enqueues by Region" on page 158 "PIMAX - Maximum Program Isolation Enqueues by Region" on page 159 "PIPL - Program Isolation Pool Percentage of Utilization" on page 159 "POOLA - Pool Allocated Storage" on page 160 "POOLN - Net Expansion Count" on page 161 "POOLT - Total Expansion/Compression Count" on page 162 "PSBP - PSB Pool Percentage of Utilization" on page 163 "PSBW - PSB Work Area Pool Percentage of Utilization" on page 163 "RECA - RECA Pool Percentage of Utilization" on page 164 "WADIO - WADS I/O" on page 164 "WKAP - General Work Area Pool Percentage of Utilization" on page 164

Table 2. IMS Activities and Resources Measured (continued)

IMS Activity/ Resource Area	**Analyzer Displays (Analyzers Reference Manual)**	**Monitors and Traces (all references in this manual)**
IMS and OS/390 interactions (IMVS area)	"DSPST - Dispatcher Statistics" "RS - Real Storage"	"CSAFR - CSA Fragmentation" on page 165 "CSAUT - CSA Percentage of Utilization" on page 165 "DLIO - DL/I EXCP Count by ddname" on page 166 "DPAGE - Demand Paging by Region" on page 166 "ECSAU - Extended CSA Percent Utilization" on page 167 "PAGE - Paging (Region)" on page 167 "SYSIO - EXCP Count by ddname" on page 168
IRLM functions (LOCK area)	"IRLM - IRLM IMS Status" "IRLMG - IRLM GLOBAL STATUS" "LCRES - IRLM Lock Contention by Resource" "LCUSR - IRLM Lock Contention by User" "LHRES - IRLM Locks Held by Resources" "LHUSR - IRLM Locks Held by User" "LUSRD - IRLM Lock User Detail"	"LDLCK - Number of Deadlocks" on page 169 "LHELD - Number of Locks Held" on page 169 "LKMAX - Maximum Locks Held by Region" on page 170 "LKREQ - Number of Lock Requests" on page 171 "LSUSP - Number of Suspensions" on page 171 "LWAIT - Region in IRLM Suspend" on page 172 "LWNUM - Number of Suspended IRLM Requests" on page 173 "PTBLK - Number of PTB Locks" on page 173 "VSEND - Number of VTAM Sends" on page 174
IMS Workload (IWKLD area)	"ISTAT - Terminal Input Status Display" chapter "OSTAT - Terminal Output Status Display" chapter "TRANQ - Transaction Queue Status Display" chapter "USER - User Status Summary" chapter "Requesting Workload Wait Data Collection (MWAIT)" chapter "Requesting Workload Wait Data Display (DWAIT)" chapter "DWAIT - Workload Wait Display" chapter	Chapter 27, "Requesting Workload Trace Data Collection (MTRAC)" on page 203 Chapter 34, "LTRAC - List of Trace Entries" on page 265 Chapter 35, "STRAC - Summary Trace Data Display" on page 273 Chapter 36, "DTRAC - Detail Trace Data Display" on page 287

Table 2. IMS Activities and Resources Measured (continued)

IMS Activity/ Resource Area	**Analyzer Displays (Analyzers Reference Manual)**	**Monitors and Traces (all references in this manual)**
IMS workload - DB2 activity (IWDB2 area)		"#CDB2 - DB2 Data Access Calls" on page 179 "#SDB2 - DB2 Nondata Access Calls" on page 180 "@PDB2 - DB2 CPU Time" on page 181
IMS workload - Fast Path activity (IWFP area)		"#CIC - Control Interval Contentions" on page 183 "#OBAW - OBA Latch Waits" on page 184 "@OBA - Overflow Buffer Usage" on page 185
IMS workload - global IMS region calls (IWGBL area)		"$CBMP - BMP and JBP Region Calls" on page 187 "$CDBT - DBCTL Region DLI Calls" on page 188 "$CMPP - MPP and JMP Region Calls" on page 189 "$CTOT - All Region Calls" on page 190
IMS workload transactions (IWTRN area)		"@ELAP - Average Elapsed Time" on page 191 "@INPQ - Average Input Queue Time" on page 193 "@RESP - Average Response Time" on page 194 "@TRSP - Average Terminal Response Time" on page 195 "#PROC - Number of Transactions Processed" on page 197

Chapter 5. Types of Services

The MVIMS services include:

Analyzers	Show formatted full-screen displays of target system status and activity.
Monitors	Measure and collect data about resource or workload performance over time and detect exception conditions.
Workload Wait	Measures and collects data about IMS workload events over time. Workload wait requires monitor data collection and analyzer display. MWAIT is a monitor service that activates workload wait data collection. DWAIT is an analyzer display of data collected by MWAIT.
Workload Trace	Tracks transaction processing through IMS. A trace requires monitor data collection and analyzer display. MTRAC is a monitor service that activates workload trace data collection. LTRAC, STRAC, and DTRAC are analyzer display services that show different views of data collected by MTRAC, from summarized statistics for all transactions to detailed events and data for one transaction.

Except when thresholds are exceeded, monitors execute independently in the background, transparent to the operator. When warnings are issued, the historical data collected by Workload Monitor and other products can be examined to determine problem causes. Workload Analyzer can be used to determine the cause of the IMS workload problem; then Resource Analyzer can be used to determine which resource is affected by the workload problem. MAINVIEW AutoOPERATOR for IMS can be used to automatically take an action when a warning is issued.

Resource Monitor has services that monitor resource usage, such as input queue length. When used together, Workload Monitor and Resource Monitor provide an early warning system for the entire IMS environment.

Requests

Services are activated as user requests.

- Analyzers can be selected from an analyzer service list or requested from a service display panel by overtyping the service name.
- Timer-driven monitors that measure resources, workloads, or wait events or track transaction processing are SET service requests.

  They can be selected from a monitor service list and started from a data entry panel. The collected data can be viewed with monitor display services similar to analyzers.
- Timer-controlled requests that activate transaction tracing can be started from a data entry panel accessed from a list of current traces.
- Timer-controlled requests for automatic analyzer display logging can be started from a data entry panel accessed from an analyzer service list.

SET requests for timer-driven services with options as keyword parameters can be generated automatically from data entry panels. Optional service parameters narrow the scope of a single request and let several requests for the same service be active concurrently. Additional requests can be made at any time, or requests can be modified or purged.

A standard set of timer-controlled monitor, workload wait, workload trace, or logging requests can be started automatically with the BBI-SS PAS.

Analyzers

The IMS analyzers are services that provide a formatted, full-screen display of IMS resource performance at your terminal in response to a request. A request for a display is made with a service select code and optional parameter. You can make the request from a service display panel or by selecting the service from an analyzer service list. The results of the service analysis are shown at your terminal and can be automatically logged to the BBI-SS PAS Image log at time-driven intervals with the SET facility.

Analyzer services that show lists of resources or workloads are scrollable.

Some analyzers display data collected by a previously started data collection service request. For example, timer-driven monitors collect short-term history data that can be displayed with a plot display. The workload wait (MWAIT and DWAIT) and trace services (MTRAC and LTRAC, STRAC, or DTRAC) in IMS Workload Analyzer function similarly. The MWAIT and MTRAC services are started like a monitor to collect data:

- MWAIT data collection can then be displayed by the DWAIT workload wait display.
- MTRAC data collection can then be displayed by:

  LTRAC — Shows a list of all traced events for an MTRAC request.

  STRAC — Shows summarized trace statistics for one transaction.

  DTRAC — Shows a chronological detail trace of events for one transaction and includes associated database I/O data and segment search argument, key feedback, and I/O area data. If a transaction has DB2, CICS, or MQSeries events associated with it, those events are also displayed.

Monitors

The IMS monitors are services that collect data measurements about resource usage or workload performance and detect warning conditions in response to a user request. A request to activate a timer-driven monitor is made by specifying a service select code and optional parameter with the SET facility. You can make a request from a service display panel or from a monitor service list. Selecting a service from a monitor list displays a data entry panel primed with the SET request keywords for the selected service. The request specifies the resource to be measured, the sampling frequency, the time of day to begin the sampling, the sampling duration, the disposition of data when the IMS subsystem is not available, and the threshold that defines a warning condition for the monitor request.

Each monitor service is a timer-driven monitor that:

- Measures resource or workload status, count of activity, or usage percentage
- Calculates a ratio of resource usage over time. The measurement obtained at each sampling interval is compared to the user-defined threshold for that request. The threshold comparison detects conditions for which user-defined warning messages can be issued. The measurements are stored online so that a plot of the recent history of a monitor request can be viewed at any time. Optional service parameters narrow the scope of a single request and let several requests for the same monitor service be active concurrently.

Starting and Stopping Monitors

With a monitor request, the user selects an IMS variable to be monitored, defines a sampling frequency, and assigns a warning threshold appropriate to the site's environment. A standard set of monitor requests can be started automatically with the BBI-SS PAS. Additional monitor requests can be made at any time or requests can be modified or purged

Monitor Response to Target IMS Shutdown and Startup

Active monitors detect target system shutdown and startup. When a request is made, quiesce and restart options can be defined that specify the monitor action when the target system stops or restarts. Monitors can:

- Be restarted automatically

  Previously collected data can be saved or deleted.
- Be purged when the target IMS stops
- Remain in a quiesced state when the target IMS starts

Data Collection

A monitor service request is activated through a timer facility, which is controlled through SET service requests. A request for a service specifies:

- The monitor to be used (service select code and parameter). This defines the IMS resource or activity to be measured.
- When monitoring should begin.
- How long monitoring should continue.
- How often the activity is to be measured (sampling interval).
- Which service functions should be performed.

The active service request automatically measures the corresponding system variable (see “Data Types Measured” on page 52) at the time interval specified on the SET request. This data can be displayed online with a plot service request. The requested plot display can be automatically logged or refreshed.

The workload monitors and resource monitors use the same timer and SET facility to activate and control requests. However, they do not use the same facility for data collection.

- Workload trace and workload monitors use the Event Collector to collect transaction-level statistics.
- Resource monitors collect data through their own timer-driven services.

Historical Data Stored

To make concise short- and long-term histories available for the graphic plot display, historical data is always stored as:

- Ten detail measurements. These are the latest 10 values, each collected at the expiration of the user-defined interval. For example, if the standard sampling interval of 1 minute is in effect, the measurements of each of the last 10 minutes are available. When a new measurement is made, the oldest value is overwritten (in other words, the values wrap around every 10 intervals).
- Two summary periods, current and previous. Both values are updated at the expiration of 10 intervals (wrap point) when 10 new detail measurements have been collected.
  The current period value is moved to the previous period and the sum of the 10 detail measurements is moved to the current period. The current period value includes from 1 to 10 of the detail measurements available at any one time. The wrap point is indicated by an arrow in the display provided by the PLOT service.
- Total. This value is the total accumulated in the time the request has been active. It is updated at each interval when a new measurement is made. Both the summary periods and the total are shown as averages-per-interval in the graphic display so the plotted detail values can be compared.
- A frequency distribution. This distribution is updated at each interval if range limits are defined with the request. From two to five ranges are allowed. The new measurement value is compared to the defined limits to find the range in which it belongs and the number of occurrences for that range is incremented by one.
- The high-water mark. This value is the maximum value ever measured at any interval and the time it occurred.

Following is an example of the storage of historical data:

If a request is started at 10:00 a.m. with an interval of one minute, the detail counters wrap around at 10:10, 10:20, and so on. The history available at 10:35 is:

- The detail measurements cover the last 10 minutes, from 10:25 to 10:35.
- The current period is from 10:20 to 10:30.
- The previous period is from 10:10 to 10:20.
- The total is from 10:00 to 10:35.

Data Types Measured

The following four types of automatic service measurements can be taken periodically and shown by the general Performance Management PLOT display request.

COUNT

An activity count over time; for example, the number of lock requests within a specified time interval.

Note: When COUNT data is plotted, in addition to the counts, rates-per-second are automatically calculated and shown for the displayed time intervals (AVG/SEC).

AVERAGE

The quantity over time; for example, the average value of transaction response time in the specified interval.

Note: When AVERAGE data is plotted, in addition to the averages, the event counts used to calculate the averages are also shown for the displayed time intervals (EVENTS); for example, for the plotted average response time, the number of transactions measured is shown.

STATUS

The status level at the moment of measurement; for example, queue lengths or level of PI enqueues.

PERCENT

Resource usage at the moment of measurement, expressed as a percentage of the maximum; for example, percent pool utilization.

See “Monitor History Display (PLOT)” on page 77 for a sample PLOT display and a complete description of the display contents.

A fifth type of data measurement, WARNING ONLY, does not collect historical data or produce a plot:

WARNING ONLY

A condition measurement that can be checked against a warning threshold; for example, the number of DL/I calls performed in a region since the last program scheduling.

Warning Conditions

SET request parameters for a monitor service can be used to define a value that is compared to the measurement taken during the requested sampling interval. The comparison establishes a warning condition when the measurement either exceeds a maximum threshold or is less than a minimum threshold. When the service detects this exception, it automatically sends a message to the BBI-SS PAS Journal and also to the system console upon user request. The message text is:

- A unique message ID
- The title of the service, which can be changed by the user with a TITLE parameter in the service request
- The measured value
- The sampling interval (if applicable)
- The defined threshold

The service sends an exception-cleared message with the service title to the BBI-SS PAS Journal when the service no longer detects a condition greater than the defined threshold value.

The monitor request can be used to:

- Specify a user-defined threshold value (WMAX or WVAL keyword).

  Each new measurement of the system variable made at the expiration of an interval is compared to the threshold value.

  When the measurement either exceeds a maximum threshold or is less than a minimum threshold, a warning condition exists and warning messages are sent automatically to the BBI-SS PAS Journal log.

  Note: The LOG DISPLAY option on the Primary Option Menu can be used to view the BBI-SS PAS Journal log.

- Send warning messages also to the OS/390 console through the write-to-operator (WTO) facility (WMSG keyword).
- Specify the number of warnings to be sent for one exception condition (WLIM keyword), the number of times the exception is detected before the first message (WIF keyword), and the number of times the exception is detected between messages (WIN keyword).

These options can be used to avoid flip-flop situations where a condition often varies just above and below the threshold, triggering many messages. For example, a condition could be checked every 30 seconds with a warning only if that condition persists for 3 minutes (WIF=6), repeated warnings only after another 5 minutes (WIN=10), and a limit of 10 (WLIM=10) warnings (the problem is known and investigated by then).

Monitor Request Title

A monitor request is identified with a title. If a parameter is specified for a requested service, the parameter is shown with the title. The title and applicable parameter for each request are shown in the:

- PLOT graphic display of the data collected by the requested monitor
- Active Timer Requests application (see "Request Status" on page 60)
- DMON or DWARN active monitor summary display
- Warning message issued when the measurement of the resource exceeds a threshold defined for the monitor

If a parameter is not used for the monitor request, the PLOT display and the warning message show the default (TOTAL) for that field.

Each service has a default title that can be customized. This process is described in "Service Table Definition" in the *MVIMS Customization Guide*. The default titles of all the monitor services are in the service descriptions in this manual.

When requesting a monitor service, you can use the TITLE keyword to make the title more meaningful to the installation for that specific request. A user-defined title can be 1 to 24 characters long. The title can be defined in a data entry panel or with a SET request in the Service Display panel. If a title is defined with a SET request, the title must be enclosed in single quotes.

Warning Message Format

Each service has a unique warning message associated with it. A warning message is issued when the condition established by the user is detected by the monitor service, as described in the preceding section. The format of a warning message is:

```
ccnnn0W (nn) hh:mm:ss title(parm) = v [IN x intrvl] [srvdata]
(>thrshld) *****
```

where:

ccnnn0W — Is the warning message ID issued by the requested monitor service.

cc	Is a two-character code indicating the service type; for example, RM indicates Resource Monitor.
nnn	Is the numerical message identifier associated with the requested service.
0	Indicates that the detected threshold currently exists.
W	Represents a warning message.

(nn) — Is the number of times the warning message was issued.

hh:mm:ss — Is a timestamp in hours, minutes, and seconds.

title — Is a default or user-defined title for the service (see "Monitor Request Title" on page 54).

(parm) Is an optional parameter that is part of the reqid (service select code plus parameter) for the SET request, as described previously. (TOTAL) is the default if a parameter is not specified for the service request.

v Is the current measured value.

IN x intrvl Is the time specified for the resource sampling with the INTERVAL keyword of the SET request where:

x Can be nn, nnnn, hh:mm:ss, or mm:ss (n is a numeric value; hh is the number of hours; mm is the number of minutes; and ss is the number of seconds).

intrvl Is units of time measurement which can be SEC or MIN.

Note: Time measurement units are not used for hh:mm:ss.

This measurement is included in the warning message only when a COUNT data type is measured (see “Data Types Measured” on page 52).

(>thrshld) Is the threshold value specified by the WMAX or WVAL keyword of the SET request for the monitor service. A < character indicates the sampled value is less than or equal to the threshold as specified by the request.

***** Is used to emphasize the message in the BBI-SS PAS Journal log.

When a threshold is exceeded, a warning message is sent and a warning condition exists. The DWARN service can be used to show all current warning conditions. When the condition that caused the warning no longer exists, the following message is issued:

```
ccnnn1I hh:mm:ss title(parm) NO LONGER > value
```

where:

ccnnn1I Is the same as the warning message ID number except a 1 replaces the 0 in the last digit and I replaces W. The 1 indicates the detected threshold no longer exists. The I indicates this message is informational.

title Is the same as the warning message title.

(parm) Is the same as the warning message parameter.

value Is the threshold value specified by the WVAL keyword of the SET request for the monitor service.

For example, if the user request is:

```
              LKREQ - IRLM LOCK REQUESTS

WVAL     ==> 5
WMSG     ==> WTO
INTERVAL ==> 00:01:00
TITLE    ==> IRLM LOCK REQUESTS
```

and the sampled measurement is greater than 5 threads at 1:00 pm, the following RM0840W message is issued:

```
RM0840W(01) 13:00:01 IRLM LOCK REQUESTS(TOTAL) = 7 (>5)
```

When the condition no longer exists, the following RM0841I message is issued:

```
RM0841I(01) 13:31:00 IRLM LOCK REQUESTS(TOTAL) NO LONGER > 5
```

Note: The target system is identified in these messages. In the Journal log, the target name is in the origin identifier field (scroll left to view). WTO messages have both the BBI-SS PAS ID and the target (`TGT=xxxx`) appended at the end of the message text.

IMS Monitor Data Display Services

The data collected by monitor service requests can be displayed online or the displays can be logged for later analysis. They can be requested from an active timer request list (see Chapter 6, “Displaying a List of Active Timer Requests (Primary Menu Option 2)” on page 65).

The monitor-collected data can be displayed by the following services:

- PLOT

  The PLOT service provides a graphic display of the history data collected for one monitor request. A range distribution of the measured values, the maximum value ever measured, and rates-per-second present effective IMS problem analysis.

- DMON

  The DMON service provides a scrollable display of the most current, active monitor measurements. Each line has a simple graphic representation of how close the measurement is to the defined warning threshold.

  DMON identifies potential problems quickly by showing several measurements together, such as the number of transactions processed, the number of database I/Os, the level of PI enqueues, and average response time. You also can use this service to see how many monitors are close to their warning thresholds.

- DWARN

  The DWARN service provides a scrollable display of current, measured values like DMON, but only for monitors with a warning condition resulting from measured values exceeding user-defined thresholds.

These displays are described in detail in Chapter 7, “Monitor Display Commands” on page 77.

Workload Wait Events

The workload wait services collect and display workload wait events for all or specific IMS workloads.

Workload wait events are sampled by a request for the MWAIT monitor. The request can be tailored to select specific components of the IMS workload for accumulated wait time. For example, the only workload components that use an IMS region are scheduling, application program, and sync point. Input and output communication and queuing events are ignored when a REGION parameter is specified.

The wait data accumulated by MWAIT is viewed by a request for the DWAIT workload analyzer. DWAIT shows the workload events that account for the IMS response time. DWAIT shows wait events by the following transaction processing event components:

- Input Communications
- Input Queue
- Scheduling
- Application Program
- Sync Point
- Output Queue
- Output Communications

All or specific transaction processing event components can be viewed.

Workload Trace

The workload trace services collect and display trace data about transaction processing and allow traced data to be logged to external VSAM data sets called trace log data sets (TLDS). Active traces can be viewed online with the current traces option on the Primary Option Menu. Logged traces can be viewed online with the history traces option on the Primary Option Menu.

The Event Collector must be active to implement a trace. As described in Chapter 27, “Requesting Workload Trace Data Collection (MTRAC)” on page 203, a trace is implemented by a request for the MTRAC monitor. The request can be tailored so that only the trace data needed to detect and solve problems is collected. The request can specify either a summary or detail trace.

The traced data collected by an MTRAC request can be viewed by requesting the LTRAC, STRAC, or DTRAC services. Display of a summary trace (LTRAC or STRAC) provides high-level quick answers about a transaction as it flows through IMS. Display of a detail trace (DTRAC) provides a chronological replay of the exact sequence of traced transaction events and includes associated database I/O data and segment search argument, key feedback, and I/O area data. If a transaction has DB2, CICS, or MQSeries events associated with it, those events are also displayed.

Logging a Display

A display can be logged in three different data sets by:

- Entering a Y for the yes option in the LOG field of the display, which records the display in the TS Image log.
- Pressing the SCREEN PRINT key (PF4/16) after the display is shown, which records the display to a BBISPRNT data set.
- Requesting automatic logging of an analyzer or monitor display service, which records the display automatically, without user interaction, to the BBI-SS PAS Image log.

DMON and DWARN display logging can also be requested for active monitors as described in "Automatic BBI-SS PAS Image Logging of Monitor Summary Displays" on page 59.

The log records can be printed offline using the IMRPRINT utility (BBSAMP member ILOGJCL), as described in the *MAINVIEW Administration Guide*.

Automatic BBI-SS PAS Image Logging of Analyzer Displays

An Image log request can be made by selecting a data entry panel from an analyzer service list with the I line command. The SET timer facility invokes an IMS analyzer service and logs the display automatically to the BBI-SS PAS Image log at a user-specified interval. `ATINTVL` is the default. For example:

```
SERV ==> SET
PARM ==> REQ=PI,I=00:05:00
```

requests Image logging of the program isolation display every five minutes. For a description of the keywords used to activate an Image log request either with a SET request or from the Image log request data entry panel.

Automatic BBI-SS PAS Image Logging of PLOT Display

Logging of a monitor PLOT display to the BBI-SS PAS Image log is coordinated automatically by specifying a LOG parameter with a monitor request (as described on 95 and page 100) or with a SET request as follows:

```
SERV ==>SET
PARM ==>REQ=HPSTL,1,I=00:01:00,LOG=ATPD
```

This request invokes the HPSTL data collection monitor to collect the number of unsuccessful Hiperspace reads for VSAM subpool 1 at one minute intervals. At the end of each complete period (`LOG=ATPD`), which is 10 intervals, a plot display of the data is logged to the BBI-SS PAS Image log. For this request, a plot is logged every 10 minutes.

A convenient logging frequency for a complete monitor history is at the end of each period (ATPD). A period is the completion of 10 time intervals. A PLOT display can also be logged at each interval (`LOG=ATINTVL`), only once at the completion of the request (`LOG=ATSTOP`), or only when a warning condition is detected (`LOG=ATWARN`).

Automatic BBI-SS PAS Image Logging of Monitor Summary Displays

A summary of active monitor status can be logged to the Image log with DMON or DWARN, as shown by the following requests which can be made from the SERV field of any display.

For example, this request logs the DMON service display, which shows the current status of the first 15 active monitors, every 10 minutes:

```
SERV ==> SET
PARM ==> REQ=DMON,I=00:10:00
```

The following request logs the DWARN service display, which shows the current status of the first 15 active monitors with a warning condition, every minute:

```
SERV ==> SET
PARM ==> REQ=DWARN,I=00:01:00
```

Table 16 on page 308 describes the keywords used to request Image logging of a DMON or DWARN display.

Note: If a parameter is not specified for a DMON or DWARN Image log request, the first 15 active monitors are logged. To log the next set of 15 active monitors, specify 16 in the PARM field, and so on.

Request Status

The status of timer-driven requests can be displayed by accessing the Active Timer Requests application:

> Select Primary Menu Option 2, MONITORS - Early Warnings/Recent History (Active Timer Requests)

The Active Timer Requests application lists standard monitor service requests and any additional timer-driven data collection services and timer-driven Image logging requests. It shows how many requests are active already and provides direct access to the data collected by that request. You can access data entry panels that allow current options to be viewed or modified, purge an active request, or use a request as a model to start a new request.

The Display Statistics and Defaults application, Primary Menu Option 5 from the Primary Option Menu, provides general information about the Timer Facility. It shows status information, some statistics, default parameters in effect, and a summary of the active timer requests for the BBI-SS PAS associated with the specified target `(TGT===>)`.

Grouping Requests

Multiple timer-driven services can be started by defining a series of requests in a member of the BBI-SS PAS BBPARM data set. The member can be started from a TS or automatically when the BBI-SS PAS starts, as described in “Request Initiation” on page 301.

BBPARM member, BLKIMFW, contains a sample starter set of IMS Resource Monitor and IMS Workload Monitor requests. Many of the requests have suggested warning thresholds; some only show activity in the IMS target.

This member should be used for automatic monitor startup (see BBPARM member BBIISP00) until a set of monitors can be customized for each IMS target at your site.

Part 3. Monitors

This section summarizes how to request a service and describes what each service does. The service descriptions are organized into groups that parallel the transaction processing sequence within IMS.

Chapter 6. Displaying a List of Active Timer Requests (Primary Menu Option 2)

The Active Timer Requests list can be accessed by using the D line command from the Data Collection Monitors list application (see Chapter 8, “Displaying a List of Data Collection Monitors (SM Command)” on page 87) or by selecting Primary Menu Option 2:

- Using the D line command displays the active timer requests for the selected service only.
- Selecting the Active Timer Requests list as a menu option displays all the active timer requests (monitor service requests, workload wait service requests, and Image log requests for analyzer or monitor summary service displays) that you are authorized to view.

```
BMC SOFTWARE  -------------- ACTIVE TIMER REQUESTS --------- COMMAND(S) ISSUED
COMMAND ===>                                                   TGT ===> IMSxxx

                                   INPUT  INTVL ==> 3        TIME --  14:32:27
 COMMANDS: SM (START MONITORS), SORT, AREA, X ON/OFF, DM (DMON), DW (DWARN)
 LC CMDS:  S (SELECT), W (SHOW),      M (MODIFY)
           P (PURGE),  R (REPLICATE), H (HELP), Z (STOP)                    >>>
 LC   SERV  PARM     TITLE                 CURRENT  WVAL  -8-6-4-2-0+2+4+6+8+
      DBIO           DB I/O COUNT BY SUBPOOL     36    10| *********W********* |
      DBIO  1        DB I/O COUNT BY SUBPOOL      0      |                     |
      DBIO  3        DB I/O COUNT BY SUBPOOL     13      |                     |
      DBIO  2        DB I/O COUNT BY SUBPOOL      0      |                     |
      DBHIT          OSAM HIT RATIO BY SP        93   <80| *******  W          |
      DBHIT 1        OSAM HIT RATIO BY SP         0   <90|          W          |
      DBHIT 3        OSAM HIT RATIO BY SP       100   <80| ******   W          |
      VHIT           VSAM HIT RATIO BY SP       100      |                     |
      OBUFW          OLDS BUFFER WAITS            0     1|          W          |
      OCHKW          OLDS CHECK WRITES           13     1| *********W********* |
      WADIO          WADS I/O                    15     1| *********W********* |
      DPAGE          DEMAND PAGING BY REGION    498   500| *********W          |
      CSAUT          CSA % UTILIZATION           56    80| ******   W          |
      ECSAU          ECSA % UTILIZATION          61    60| *********W          |
      DLIO  ++++++++ DL/1 EXCP COUNT BY DDNAM    31      |                     |
      SYSIO ++++++++ EXCP COUNT BY DDNAME        21      |                     |
      SHMSG          SHORT MSG QUEUE % UTIL       0    10|          W          |
      LGMSG          LONG MSG QUEUE % UTIL        4    10| ***      W          |
      QBLKS          QBLKS % UTILIZATION          0    10|          W          |
      LHELD          IRLM LOCKS HELD              2      |                     |
      LKREQ          IRLM LOCK REQUESTS           0      |                     |
      LSUSP          IRLM SUSPENSIONS             0      |                     |
       **************************** END OF REQUESTS ****************************
```

Figure 1. List Active Timer Requests Application (Before Scrolling Right)

```
BMC SOFTWARE  -------------- ACTIVE TIMER REQUESTS ---------- FUNCTION COMPLETE
  COMMAND ===>        TGT ===> IMSxxx
                                     INPUT    INTVL ==> 3      TIME --  07:46:04
  COMMANDS: SM (START MONITORS), SORT, AREA, X ON|OFF, DM (DMON), DW (DWARN)
  LC CMDS:  S (SELECT), W (SHOW),      M (MODIFY),
            P (PURGE),  R (REPLICATE), H (HELP),   Z (STOP),                 <<<
  LC   SERV  PARM     TITLE                      USER ID  TARGET  SEC AREA  STAT
       DBIO           DB I/O COUNT BY SUBPOOL    LAA1     IMSxxx   A  DB    ACTV
       DBIO  1        DB I/O COUNT BY SUBPOOL    LAA1     IMSxxx   A  DB    ACTV
       DBIO  3        DB I/O COUNT BY SUBPOOL    LAA1     IMSxxx   A  DB    ACTV
       DBIO  2        DB I/O COUNT BY SUBPOOL    LAA1     IMSxxx   A  DB    ACTV
       DBHIT          OSAM HIT RATIO BY SP       LAA1     IMSxxx   A  DB    ACTV
       DBHIT 1        OSAM HIT RATIO BY SP       LAA1     IMSxxx   A  DB    ACTV
       DBHIT 3        OSAM HIT RATIO BY SP       LAA1     IMSxxx   A  DB    ACTV
       VHIT           VSAM HIT RATIO BY SP       LAA1     IMSxxx   A  DB    ACTV
       OBUFW          OLDS BUFFER WAITS          LAA1     IMSxxx   A  INTNL ACTV
       OCHKW          OLDS CHECK WRITES          LAA1     IMSxxx   A  INTNL ACTV
       WADIO          WADS I/O                   LAA1     IMSxxx   A  INTNL ACTV
       DPAGE          DEMAND PAGING BY REGION    LAA1     IMSxxx   A  IMVS  ACTV
       CSAUT          CSA % UTILIZATION          LAA1     IMSxxx   A  IMVS  ACTV
       ECSAU          ECSA % UTILIZATION         LAA1     IMSxxx   A  IMVS  ACTV
       DLIO  ++++++++ DL/1 EXCP COUNT BY DDNAM   LAA1     IMSxxx   A  IMVS  ACTV
       SYSIO ++++++++ EXCP COUNT BY DDNAME       LAA1     IMSxxx   A  IMVS  ACTV
       SHMSG          SHORT MSG QUEUE % UTIL     LAA1     IMSxxx   A  QUEUE ACTV
       LGMSG          LONG MSG QUEUE % UTIL      LAA1     IMSxxx   A  QUEUE ACTV
       QBLKS          QBLKS % UTILIZATION        LAA1     IMSxxx   A  QUEUE ACTV
       LHELD          IRLM LOCKS HELD            LAA1     IMSxxx   A  LOCK  ACTV
       LKREQ          IRLM LOCK REQUESTS         LAA1     IMSxxx   A  LOCK  ACTV
       LSUSP          IRLM SUSPENSIONS           LAA1     IMSxxx   A  LOCK  ACTV
        **************************** END OF REQUESTS ****************************
```

Figure 2. List Active Timer Requests Application (After Scrolling Right)

A select line command for an active monitor request provides direct access to a plot display of data collected by the selected monitor. Other line commands can be used to view current timer request options, access data entry panels to replicate or change the current options to make another unique SET timer request for that service, or confirm a purge of the selected request.

Timer requests are shown in the Active Timer Requests list as they are processed. This list can be scrolled left and right and up and down. It allows service selection by line command and shows all the active requests, parameters specified for each requested service, the latest measured value, the specified warning threshold, a plot for the current sampling, the user logon identification, the target IMS of the request, the service security classification, the area of IMS being monitored, and the service status by:

Field Name	Description
INTVL	Screen refresh field. You can specify a refresh interval for this application. Entering GO in the COMMAND field or pressing PF6/18 starts the refresh cycle.
LC	A line command input field. One-character line commands can be entered in this field to view, modify, or replicate the options for a selected request; purge a request; or display HELP information about the service (see “Line Commands” on page 71).

SERV
: A scrollable list of services for all active timer requests by service select code. The Active Timer Requests list application also can be displayed by using the D line command from the Data Collection Monitors list, as described in "Line Commands" on page 90. Only those requests that are active for that service are shown.

PARM
: This field shows the parameters that were defined for the active requests.

TITLE
: The service title.

CURRENT
: The latest measured value.

  Note: If the request is not active, its status (as defined in the STAT field) is displayed in this column.

WVAL
: The warning threshold.

`-8-6-4-2-0+2+4+6+8+0`
: A plot for the current sampling and a warning threshold (W marker) if the WVAL keyword was specified with the SET request. Plot characters indicate a trend:

  - < Shows a downward trend from the preceding sampled values.
  - > Shows an upward trend from the preceding sampled values.
  - * Shows no change from the preceding sampled values.

  If you have a color monitor, the graph is displayed in these colors:

Red	Warning status
Turquoise	Normal values for the current interval
Yellow	Maximum thresholds: values for the current period are greater than the values for the previous period Minimum thresholds: values for the current period are less than the values for the previous period

USER ID
: The logon identification of the user who made the request.

TARGET
: The IMS subsystem defined as the target of the requested service either by default or user-specified.

SEC
: The security code for user access to the service.

AREA
: The IMS resource area being analyzed. This field could contain:

Field Data	Description
DB	Database activity and buffer pool utilization
IMVS	IMS and OS/390 interactions
INTNL	IMS internal functions
IWDB2	IMS workload DB2 activity
IWFP	IMS workload Fast Path activity
IWGBL	IMS workload global IMS region calls

IWKLD	IMS workload
IWTRN	IMS workload transactions
LOCK	RLM functions
MFS	Terminal I/O
QUEUE	IMS queuing
REGN	Application program activity in the dependent regions
SCHED	Scheduling of application programs in the dependent region

STAT The service request status which could be:

Field Data	Description
ACTV	The request is active.
COMP	The request executed and completed normally.
HELD	The request is being held and is pending release.
INIT	The request is being invoked for the first time (a start time was specified, but it has not been reached).
INV	The request terminated because of an invalid parameter or measurement. The BBI-SS PAS Journal log contains a descriptive message of the error.
LOCK	A LOCK command was issued for the service or the service abended.
QIS	The service is quiesced, because the target IMS is not active.
RST	The target IMS restarted. The request is waiting until the current interval expires before performing restart processing as specified by the RST keyword in the original request.

Application Transfer Commands

The following related application transfer commands can be entered in the COMMAND field of the Active Timer Requests list application showing all timer requests:

AN (Analyzer Display Services)
AT (Active Timer Requests)
CT | VT (View Current Traces)
DM (DMON)
DW (DWARN)
HT (History Traces)
SM (Start Monitors)
ST (Start Traces)

AN (Analyzer Display Services)

AN displays the analyzer display services application, which lists all the analyzers that you are authorized to view and allows selection of analyzers with the S line command. The application shows allowable parameters for each service, the service security classification, the area of IMS being analyzed, and the service status.

AT (Active Timer Requests)

AT displays the active timer requests application, which lists all the monitors that have been activated and allows selection of monitors with the S line command. Activated monitors are started either manually (with the SM command) or automatically (by a BLK member defined to start automatically using the BBPARM member BBIISP00). The application also shows workload wait service requests and image log requests for analyzers.

CT | VT (View Current Traces)

CT displays the Current Traces application, which lists all currently active traces. From this application, you can use the S line command to display collected trace data.

DM (DMON)

DM displays the DMON service (see "Active Monitor Summary Display (DMON)" on page 83), which displays global data about the status and operation of active monitor services. If a warning threshold (WVAL) is specified with the monitor request, a graphic trend display is shown.

DW (DWARN)

DW displays the DWARN service (see "Active Monitor Warning Display (DWARN)" on page 85), which displays the status of active monitor services that are currently exceeding the threshold limit set with the WVAL parameter of the monitor request.

HT (History Traces)

HT displays the History Traces application, which is used to manage trace log data sets where traces are recorded (see Chapter 28, “Logging a Trace” on page 221).

SM (Start Monitors)

SM displays the Data Collection Monitors application which lists all of the data collection monitor services you are authorized to view. From this panel, you can use the S line command to select the data entry panel to start a timer request for any of the listed monitors.

ST (Start Traces)

ST displays the MTRAC data entry panel to start a trace request. Specify the keyword parameters, selection criteria, and exception filters in successive panels to activate data collection for a summary or detail trace.

SORT Primary Command

When the list of active timer requests is displayed initially, the list is sorted in the order requests are made. SORT can be entered in the COMMAND field of the display to sort the list by any of the column headings. The first two characters of the column heading are used with SORT as:

```
SORT cc
```

where `cc` can be any of the following two characters, which are described below in alphabetical order.

AR Sorts the list by the resource area (AREA column).
CU Sorts the numerical values in descending order (CURRENT column).
SC Sorts the list alphabetically by the security code.
SE Sorts the list alphabetically by service name (SERV column).
ST Sorts the list alphabetically by the service status displayed (STAT column).
TA Sorts the list alphabetically by target ID (TARGET ID column).
TI Sorts the list alphabetically by service title (TITLE column).
US Sorts the list alphabetically by user ID (USER ID column).
WV Sorts the numerical values in descending order (WVAL column).

AREA Primary Command

You can use the AREA command to list only the services related to a specific area. The possible areas that can be specified are shown in the AREA column. For example, to list only the IMS database services, type in the COMMAND field:

```
AREA DB
```

Type AREA without an area name to return to the list of all the services.

X ON | OFF Primary Command

To display only the requests that are in warning status, type X ON in the command line and press ENTER. To display all requests, type X OFF in the command line and press ENTER. The default is to display all requests.

Line Commands

Entering one of the following one-character line commands in the LC field for a service executes the line command function. Multiple selections can be entered at one time by selecting a series of services and pressing the ENTER key. Each display in a series is processed by pressing the END key. Each data entry timer request panel in a series that is to be modified or purged is submitted by pressing the ENTER key and then the END key to process the next request.

S SELECT. Displays a plot of collected data.

Selection of an active request for a resource or workload monitor service displays a plot of the data collected by the selected request (see “History PLOT Display (S Line Command for Monitor Request)”).

Note: The S line command (SELECT a plot) is not valid for an active Image log request of an analyzer or general service display.

W SHOW. Shows a display panel of the timer request options defined for the selected request for viewing only (see “Show Timer Request (W Line Command)” on page 76).

M MODIFY. Shows a data entry panel of the timer request options defined for the selected request so the options can be changed (see “Modify Timer Request (M Line Command)” on page 72).

P PURGE. Displays a PURGE panel to verify a purge of the selected request (see “Purge Request (P Line Command)” on page 74).

R REPLICATE. Shows a data entry panel of the timer request options defined for the selected request so that the options can be repeated or changed to make a new request for that service (see “Replicate Timer Request (R Line Command)” on page 75). The request must be unique (defined by the service select code plus a parameter).

H HELP. Displays HELP information about the service for the selected request.

Using H for a display service request shows the service title, describes what the service does, and defines any parameters.

Using H for a monitor request shows the service title, describes the value measured by that monitor service, defines any parameters, and shows the format of the monitor warning message.

Z STOP. Stops the request and retains collected data. The STOP time equals the current time.

History PLOT Display (S Line Command for Monitor Request)

The S line command displays a graphic history plot of the data collected by the selected monitor service. It generates the display described in “Monitor History Display (PLOT)” on page 77.

Modify Timer Request (M Line Command)

Selecting a request with the M line command displays a data entry panel with the options that were defined to activate data collection for a workload or monitor service or Image logging for a display service. Previously defined options that are prefixed by ===> can be changed, as shown by the example in Figure 3.

```
BMC SOFTWARE  --------- MODIFY RESOURCE MONITOR REQUEST -----  PERFORMANCE MGMT
COMMAND ===>                                                     TGT ===> IMSA
                      DBIO  - DB I/O COUNT BY SUBPOOL

 PARM:                                 (Resource Selection Parameter)

 INTERVAL:   00:01:00  START:   14:26:00   STOP ==>               QIS ==> YES

 WVAL    ==> 10       WMSG  ==>            WLIM ==> 10  WIF ==> 1   WIN ==> 1

 RST     ==> HOT                           (Restart Option: HOT,COLD,PUR,QIS)

 TITLE:      DB I/O COUNT BY SUBPOOL       (Title)

 PLOTMAX ==>                               (Maximum PLOT X-Axis Value)

 RANGES:                                   (1-4 Range Distr. Upper Limits)

 LOG     ==> NO                            (NO,ATSTOP,ATPD,ATINTVL,ATWARN)
```

Figure 3. Modify Timer Request Panel (M Line Command)

Options shown with a colon (:) suffix cannot be changed. The request is submitted when the ENTER key is pressed. A short message in the upper right corner of the display shows the result of the request. If an ERROR IN REQUEST message is displayed, a short explanatory message is displayed on the third line. Pressing the END key (PF3/15) redisplays the Active Timer Requests list.

Resource Monitor Request

The M line command for a resource monitor service displays the timer request options used to start data collection (see “Start Data Collection Timer Request (S Line Command)” on page 91). As shown in Figure 3 on page 72, the following options are followed by an arrow (==>), which means that you can change their displayed values:

SET Option	Description
STOP	Service stop time
QIS	Service quiesce state
WVAL	Warning threshold
WMSG	Warning message routing
WLIM	Maximum warning messages
WIF	Number of intervals before first warning
WIN	Number of intervals between warnings
RST	Service restart
PLOTMAX	Maximum X-axis value for plot display
LOG	Automatic BBI-SS PAS Image logging of PLOT display (default is NO)

Workload Monitor Request

The options that can be modified for workload monitor requests are the same as those listed for resource monitor requests.

Note: Workload monitor selection criteria cannot be modified.

Workload Wait Request

The following options can be modified for workload wait requests:

SET Option	Description
STOP	Service stop time
RST	Service restart

Workload Trace Request

The following options can be modified for workload trace requests:

SET Option	Description
STOP	Service stop time
RST	Service restart
WRAP	Trace data wrap If you request trace logging with the MTRAC request, see “Modify a Trace Log Request (M Line Command)” on page 236 for the logging options that you can modify after a trace is active.

Image Log Request

The M line command for an analyzer or monitor summary service displays the SET timer request options used to log the service display to the BBI-SS PAS Image log (see “Start Image Log Request” in Chapter 6 of the *Analyzers Reference Manual*). The following options have fields prefixed by ===> ; this means their displayed values can be changed:

SET Option	Description
STOP	Service stop time
QIS	Service quiesce state
RST	Service restart

Purge Request (P Line Command)

Selecting a service with the P line command displays a purge confirmation panel, shown by Figure 4.

```
BMC SOFTWARE  --------------- CONFIRM PURGE REQUEST ---------- PERFORMANCE MGMT
COMMAND ===>        TGT ===> IMSA


                       DBIO  - DB I/O COUNT BY SUBPOOL


 PARM:                                      Parameter / Identifier
 STATUS:        ACTV

 START:         14:26:00  (0    days)
 STOP:


   INSTRUCTIONS:

     Press ENTER key to confirm purge request.

     Enter END command to cancel purge request.
```

Figure 4. Purge Request Application (P Line Command)

Pressing the ENTER key confirms a purge of the selected service request. A short message in the upper right corner of the display shows the result of the request. If an ERROR IN REQUEST message is displayed, a short explanatory message is displayed on the third line. Pressing the END key (PF3/15) redisplays the Active Timer Requests list.

To stop a request and retain online plot or trace, use the modify command and enter a stop time. This allows information to remain after collection stops. If a request is purged, all data is lost.

Replicate Timer Request (R Line Command)

The R line command displays a data entry panel for the selected service, as shown in Figure 5.

```
BMC SOFTWARE  -------- REPLICATE RESOURCE MONITOR REQUEST ---  PERFORMANCE MGMT
COMMAND ===>        TGT ===> IMSA

                       DBIO  - DB I/O COUNT BY SUBPOOL

 PARM     ==>                             (Resource Selection Parameter)

 INTERVAL ==> 00:01:00 START ==>          STOP ==>                QIS ==> YES

 WVAL     ==> 10       WMSG  ==>          WLIM ==> 10  WIF ==> 1   WIN ==> 1

 RST      ==> HOT                         (Restart Option: HOT,COLD,PUR,QIS)

 TITLE    ==> DB I/O COUNT BY SUBPOOL     (Title)

 PLOTMAX  ==>                             (Maximum PLOT X-Axis Value)

 RANGES   ==>                             (1-4 Range Distr. Upper Limits)

 LOG      ==> NO                          (NO,ATSTOP,ATPD,ATINTVL,ATWARN)
```

Figure 5. Replicate Timer Request Panel (R Line Command)

The options that were defined to start the request are displayed (see “Start Data Collection Timer Request (S Line Command)” on page 91). All values prefixed by ===> can be modified. This application can be used to start a new request for the selected service. Each request must be unique and is defined by the service select code and an optional parameter. The display is preset with the service select code.

Note: If the request specified a STOP time, it appears as nnnn, where nnnn is the number of intervals remaining until the monitor stops.

Pressing the ENTER key submits the request. A short message in the upper right corner of the display shows the result of the request. If an ERROR IN REQUEST message is displayed, a short explanatory message is displayed on the third line. Pressing the END key (PF3/15) redisplays the Active Timer Requests list.

Monitor or Trace Request

Using the R line command for a monitor or trace service displays all of the options previously defined to start data collection for the selected request (see “Start Data Collection Timer Request (S Line Command)” on page 91). All the option values can be replicated or changed and submitted by using the ENTER key, as long as the request is unique.

Image Log Request

Using the R line command for a display service shows all of the options previously defined to log the display to the BBI-SS PAS Image log (see “Start Image Log Request” in Chapter 6 of *MAINVIEW for IMS Online – Analyzers Reference Manual*). All the option values can be replicated or changed and submitted by using the ENTER key, as long as the request is unique.

Show Timer Request (W Line Command)

The W line command can be used for any request shown in the Active Timer Requests list. The request types in the list are either for time-driven data collection monitor services, workload collection requests, or Image logging requests of monitor summary (DMON or DWARN) or analyzer service displays. Selecting a request with the W line command generates a display panel of the previously defined options for that SET timer request, as shown by the example in Figure 6.

```
BMC SOFTWARE  ----------- SHOW RESOURCE MONITOR REQUEST -----  PERFORMANCE MGMT
COMMAND ===>        TGT ===> IMSA

                        DBIO  - DB I/O COUNT BY SUBPOOL

   PARM:                                 (Resource Selection Parameter)

   INTERVAL:   00:01:00  START:   14:26:00   STOP:                   QIS:   YES

   WVAL:       10        WMSG:               WLIM:   10   WIF:   1    WIN:   1

   RST:        HOT                           (Restart Option: HOT,COLD,PUR,QIS)

   TITLE:      DB I/O COUNT BY SUBPOOL       (Title)

   PLOTMAX:                                  (Maximum PLOT X-Axis Value)

   RANGES:                                   (1-4 Range Distr. Upper Limits)

   LOG:        NO                            (NO,ATSTOP,ATPD,ATINTVL,ATWARN)
```

Figure 6. Show Timer Request Panel (W Line Command)

Each option is suffixed by a colon (:), which means the option value cannot be changed. Pressing the END key redisplays the Active Timer Requests list.

Monitor Request

Selecting a request for a monitor, workload wait, or workload trace service with the W line command shows the options that were requested to activate data collection by that service. It is used to only view the options not to change them. Figure 6 is an example of a resource monitor request.

Image Log Request

The W line command for a monitor summary (DMON or DWARN) or analyzer service display shows the previously defined options for BBI-SS PAS Image logging with a colon suffix (see “Start Image Log Request” in Chapter 6 of *MAINVIEW for IMS Online – Analyzers Reference Manual*). It is used to only view the options not to change them.

Chapter 7. Monitor Display Commands

The monitor display commands are used to activate services that display monitor-collected data. These services are:

PLOT	Graphic plot display of data collected by a requested monitor
DMON	Status display of active monitors
DWARN	Activity summary and status display of monitor warnings

Monitor History Display (PLOT)

```
BMC SOFTWARE  -------------    MONITOR HISTORY      -----------PERFORMANCE MGMT
  SERV ==> PLOT   INPUT    14:41:46  INTVL=> 3  LOG=> N  TGT==> IMSA
  PARM ==> DBIO                 SCROLL=> N/A
  TOTAL         348       DB I/O COUNT BY SUBPOOL                START  14:26:00
     15 SAMPLES      PERIOD  00:10:00   INTERVAL  00:01:00      ELAPSED  00:15:00
  AVG/INTVL--. ----+----I----+----I----+----I----+----I----+----- .--- AVG/SEC --
    TOTAL    . *********W*************                            .      .39
    PREV PD  .          W                                         .
    CURR PD  . *********W************************                 .      .57
  INTERVALS--. ----+----I----+----I----+----I----+----I----+----- .
   14:32:00  . *********W**************************               .      .60
   14:33:00  . *********W************************                 .      .57
   14:34:00  . *********W**********************                   .      .53
   14:35:00  . *********W********************                     .      .50
   14:36:00->. *********W***************************              .      .62
   14:37:00  . *******  W                                         .      .12
   14:38:00  .          W                                         .
   14:39:00  .          W                                         .
   14:40:00  .          W                                         .
   14:41:00  .          W                                         .
             . ----+----I----+----I----+----I----+----I----+----- .
             0         10        20        30        40        50 MAX AT 14:28:00
          MAX=====>40
```

Figure 7. PLOT Sample Display

Description: The PLOT service generates a graphic display of historical data collected by a monitor service. It can be requested with:

- An S line command for an active monitor from the active timer list application (Primary Option Menu 2)
- A PLOT command as shown in Figure 7.

Color: If you have a color monitor:

Red	Indicates a warning status.
Turquoise	Indicates normal values for the current interval.
Yellow	Indicates maximum thresholds: values for the current period are greater than the values for the previous period
	Indicates minimum thresholds: values for the current period are less than the values for the previous period

Select Code: PLOT

Parameter: reqid

Field Descriptions: Each of the fields is shown and described below by display area.

Area 1 - Monitor Statistics

```
TOTAL 348 DB I/O COUNT BY SUBPOOL START  14:26:00
  15 SAMPLES PERIOD  00:10:00 INTERVAL  00:01:00  ELAPSED  00:15:00
```

The field descriptions are arranged in alphabetical order.

ELAPSED The cumulative active time for the request. If the request is not currently active, this field contains the elapsed time at the point when the request became inactive.

INTERVAL The request sampling interval.

PERIOD The elapsed time for one period. It is equal to 10 times the interval value.

SAMPLES The number of data samples collected by this request.

START The request start time. The request may go inactive at start time if the IMS subsystem is not active and `QIS=YES` (the default) was specified when the service was requested.

TOTAL Displayed for count-type services only. This field shows the sum of all values collected by this request.

Area 2 - Averages Per Interval

```
AVG/INTVL--. ----+----I----+----I----+----I----+----I----+----- .--- AVG/SEC --
  TOTAL    . *********W*************    .      .39
  PREV PD  .    W                          .
  CURR PD  . *********W*************************                .      .57
```

The field descriptions are arranged in alphabetical order.

CURR PD
: The average sample value over the most recently completed period of 10 interval samples. At the end of the period, the PREV PD value is replaced with this average.

PREV PD
: The average sample value for the preceding period of 10 interval samples.

TOTAL
: The average sample value over the life of the service.

W
: The warning threshold marker (a WVAL warning threshold must be defined with the SET request for the service). If the PLOTMAX value is less than the WVAL value, the W marker is not displayed.

Area 3 - Interval Sample Values

```
INTERVALS--. ----+----I----+----I----+----I----+----I----+----- .
 14:32:00  . *********W**************************             .      .60
 14:33:00  . *********W************************               .      .57
 14:34:00  . *********W**********************                 .      .53
 14:35:00  . *********W********************                   .      .50
 14:36:00->. *********W*****************************          .      .62
 14:37:00  . *******  W                                       .      .12
 14:38:00  .          W                                       .
 14:39:00  .          W                                       .
 14:40:00  .          W                                       .
 14:41:00  .          W                                       .
           . ----+----I----+----I----+----I----+----I----+----- .
           0         10        20        30        40        50 MAX AT 14:28:00
```

Each line shows the time the sample was taken. The X-axis scale value to be plotted is automatically adjusted to the next highest multiple of 50. If PLOTMAX is specified, values exceeding this limit are expressed with an asterisk extending beyond the right side of the graph.

->	Designates that this interval sample value and all interval sample values above this line are included in the CURR PD, shown in the Area 2 sample display.
MAX AT or MIN AT	The time the maximum or minimum sample value was measured.

Area 4 - Monitor Measurements

```
For count services:      |   For average services:
                         |
       --- AVG/SEC ---   |   --- EVENTS ---
     57.21               |       13,817
     74.99               |        1,706
     42.62               |        1,520
                         |
     19.17               |          163
     15.78               |          181
     91.15               |          140
    109.57               |          132
     22.97               |          158
     29.08               |          119
     21.55               |          129
     30.32               |          132
     45.33               |          196
     75.97               |          165
```

This area is displayed for count-type or average-type services only.

For services that measure a count, the rates per second are shown for the displayed time intervals. These values are calculated by dividing the number of event occurrences by the elapsed time.

For services that measure an average, the event counts used to calculate the averages, such as the number of transactions for which response time was measured, are shown.

Area 5 - Range Limits and Distribution

```
RANGE:    0->1220  1221->2440  2441->3660  3661->4880  4881->6758 ---TOTAL---
DISTR:    2    7%     8   29%     7   25%     4   14%     7   25%    28  100%
```

This area is displayed only if the RANGES keyword is specified with the SET timer request for the service. Otherwise, only the maximum or minimum value ever measured is shown as MAX or MIN value.

The descriptions are arranged in alphabetical order.

DISTR: The number and relative percentage of sample values falling into each range.

RANGE: The user-defined range limits. Four upper-range limits can be defined with the RANGES keyword of the SET request. The first low range is set to zero. Each successive lower-range limit is equal to the preceding upper-range limit +1.

The fifth upper-range limit is the sample value that exceeded the maximum range limit specified by the RANGES keyword of the SET request. It is always the maximum observed sample value. If no sample value exceeds the maximum specified range limit, the maximum value encountered is displayed.

Active Monitor Summary Display (DMON)

```
BMC SOFTWARE  -------------    MONITOR SUMMARY      -----------PERFORMANCE MGMT
SERV ==> DMON            INPUT    14:42:30  INTVL=> 3  LOG=> N  TGT==> IMSA
PARM ==> 15             IM1714I CURRENT DISPLAY POSITION =   1.  SCROLL=> N/A
EXPAND:  LINESEL(PLOT)
---------------------------------- STATISTICS---------------------------------
STATUS  - ACTIVE        10:51:25 IMF-SS STARTED      EVENT COLLECTOR - ACTIVE
MONITOR STATUS:  26 ACTIVE     0 COMPLETE     0 QUIESCED    4 CURRENT WARNINGS
------------------------------MONITOR STATISTICS------------------------------
REQUEST                  --PARM-- CURRENT WVAL| -8-6-4-2-0+2+4+6+8+ |INTVL WARN
DB I/O COUNT BY SUBPOOL                0    10|         W          | 1M
DB I/O COUNT BY SUBPOOL  1             0      |                    | 1M
DB I/O COUNT BY SUBPOOL  3             0      |                    | 1M
DB I/O COUNT BY SUBPOOL  2             0      |                    | 1M
OSAM HIT RATIO BY SP                   0   <80|         W          | 1M     5M
OSAM HIT RATIO BY SP     1             0   <90|         W          | 1M    15M
OSAM HIT RATIO BY SP     3             0   <80|         W          | 1M     5M
VSAM HIT RATIO BY SP                   0      |                    | 1M
OLDS BUFFER WAITS                      0     1|         W          | 1M
OLDS CHECK WRITES                      0     1|         W          | 1M
WADS I/O                               0     1|         W          | 1M
DEMAND PAGING BY REGION              315   500| >>>>>    W          | 1M
CSA % UTILIZATION                     56    80| >>>>>>   W          | 1M
ECSA % UTILIZATION                    61    60| >>>>>>>>>W          | 1M    12M
```

Figure 8. DMON Sample Display

Description: The DMON service displays the current status

DMON of all active monitors. It shows each request, the current sampling, the defined threshold, and a graphic summary of all the current measured values compared to defined thresholds.

It can be requested with a:

- DM
- DMON

command as shown in Figure 8.

There is a one-line display for each request. Each line shows the:

- Full request title (REQUEST)
- Request parameter (PARM)
- Latest measured value (CURRENT)
- Warning threshold (WVAL) (minimum thresholds are prefixed by <)
- Plot for the current sampling and a warning threshold (W marker) if the WVAL keyword was specified with the SET request
- Sampling interval (INTVL)
- Duration of the current warning condition (WARN)

A -100 to +100 warning threshold percentage can be plotted. The plot characters indicate a trend:

< Shows a downward trend from the preceding sampled values

\> Shows an upward trend from the preceding sampled values

* Shows no change from the preceding sampled values

Previous period to current period (see the “Area 2” description in “Monitor History Display (PLOT)” on page 77) comparison determines the trend direction.

Color: If you have a color monitor:

Red	Indicates a warning status.
Turquoise	Indicates normal values for the current interval.
Yellow	Indicates maximum thresholds: values for the current period are greater than the values for the previous period.
	Indicates minimum thresholds: values for the current period are less than the values for the previous period.

Select Code: DMON

Parameter: Specifies the monitor requests to be displayed. The status of 15 requests are displayed. They can be selected by:

- Entering the relative request number (1 to 3 digits); the default is 1. DMON displays 15 monitor requests beginning with the selected request. Press ENTER to see the remaining currently defined monitor requests in groups of 15.
- Entering a valid reqid (see “Service Request ID (reqid)” on page 301). DMON displays 15 monitor requests beginning with the specified request. Press ENTER to see the remaining currently defined monitor requests in groups of 15.
- Entering a 0. DMON displays only the first 15 monitor requests. The parameter remains at 0 until either a relative number or reqid is entered.

Expand: The DMON display can be EXPANDed to the following display:

LINESEL(PLOT)
PLOT for a specific monitor can be selected by positioning the cursor in one of the request lines and pressing ENTER.

Active Monitor Warning Display (DWARN)

```
BMC SOFTWARE  -------------    WARNING SUMMARY      -----------PERFORMANCE MGMT
SERV ==> DWARN           INPUT    14:42:50  INTVL=> 3  LOG=> N  TGT==> IMSA
PARM ==> 1              IM1714I CURRENT DISPLAY POSITION =   1.  SCROLL=> N/A
EXPAND:  LINESEL(PLOT)
---------------------------------- STATISTICS----------------------------------
STATUS  - ACTIVE        10:51:25 IMF-SS STARTED       EVENT COLLECTOR - ACTIVE
MONITOR STATUS:  26 ACTIVE     0 COMPLETE     0 QUIESCED    4 CURRENT WARNINGS
--------------------------CURRENT WARNING CONDITIONS---------------------------
REQUEST                        --PARM-- CURRENT WVAL     0+2+4+6+8+ |INTVL WARN
DBHIT OSAM HIT RATIO BY SP                    0   <80     W          | 1M     5M
DBHIT OSAM HIT RATIO BY SP     1              0   <90     W          | 1M    15M
DBHIT OSAM HIT RATIO BY SP     3              0   <80     W          | 1M     5M
ECSAU ECSA % UTILIZATION                     61    60     W          | 1M    12M
------------------------------END REQUEST BLOCKS-------------------------------
```

Figure 9. DWARN Sample Display

Description: The DWARN service displays only those active monitors currently in a warning condition. It shows each active warning request, the current sampling, the defined threshold, and a graphic summary of the measured values that exceeded defined thresholds.

It can be requested with a:

- DW
- DWARN

command as shown in Figure 9.

There is a one-line display for each request. Each line shows the:

- Full request title (REQUEST)
- Request parameter (PARM)
- Latest measured value (CURRENT)
- Warning threshold (WVAL) (minimum thresholds start with <)
- Plot for the current sampling and a warning threshold (W marker)
- Sampling interval (INTVL)
- Duration of the current warning condition (WARN)

A -100 to +100 warning threshold percentage can be plotted. The plot characters indicate a trend:

< Shows a downward trend from the preceding sampled values

> Shows an upward trend from the preceding sampled values

* Shows no change from the preceding sampled values

Previous period to current period (see the “Area 2” description in “Monitor History Display (PLOT)” on page 77) comparison determines the trend direction.

Color: If you have a color monitor:

Red	Indicates a warning status.
Turquoise	Indicates normal values for the current interval.
Yellow	Indicates maximum thresholds: values for the current period are greater than the values for the previous period.
	Indicates minimum thresholds: values for the current period are less than the values for the previous period.

Select Code: DWARN

Parameter: Specifies the warning monitor requests to be displayed. The status of 15 requests in a warning condition are displayed. They can be selected by:

- Entering the relative request number (1 to 3 digits); the default is 1. DWARN displays 15 warning monitor requests beginning with the selected request. Press ENTER to see the remaining currently defined warning monitor requests in groups of 15.
- Entering a valid reqid (see “Service Request ID (reqid)” on page 301). DWARN displays 15 warning monitor requests beginning with the specified request. Press ENTER to see the remaining currently defined warning monitor requests in groups of 15.
- Entering a 0. DWARN displays only the first 15 warning monitor requests. The parameter remains at 0 until either a relative number or reqid is entered.

Expand: The DWARN display can be EXPANDed to the following display:

LINESEL(PLOT)
PLOT for a specific monitor can be selected by positioning the cursor in one of the request lines and pressing ENTER.

Chapter 8. Displaying a List of Data Collection Monitors (SM Command)

This application is a scrollable list of all the data collection monitor services. To view a list of all data collection monitors:

1. Select Primary Option Menu 2, MONITORS, from the Primary Option Menu for a list of active monitor requests.
2. Enter SM in the COMMAND line. A scrollable list of all the data collection monitors that you are authorized to view is displayed as shown below.

```
BMC SOFTWARE  -------------- DATA COLLECTION MONITORS -------- PERFORMANCE MGMT
COMMAND ===>
                                             TGT ===> IMSA
COMMANDS: SORT,AREA
 LC CMDS: S(SET UP), D(DISPLAY ACTIVE), H(HELP)
 LC   SERV    # ACTIVE  TITLE                    PARM TYPE    SEC  AREA   STAT
      MWAIT        1    MONITOR WORKLOAD WAIT    (IDENTIFIER)  A   IWKLD
      MTRAC        1    WORKLOAD TRACE           (IDENTIFIER)  A   IWKLD
      MFSIO             MFS I/O                                A   MFS
      MFSIR             MFS IMMEDIATE REQUESTS                 A   MFS
      MFSFD             % MFS BLKS FOUND IN POOL               A   MFS
      QIO               QUEUE I/O                              A   QUEUE
      QWAIT             QUEUING WAITS                          A   QUEUE
      INQBG             IN-Q LENGTH BY BALG      (IDENTIFIER)  A   QUEUE
      INQTR             IN-Q LENGTH BY TRANCODE  (TRANCODE)    A   QUEUE
      INQCL             IN-Q LENGTH BY CLASS     (CLASS#)      A   QUEUE
      IQSCL             SCHDULABLE IN-Q BY CLASS (CLASS#)      A   QUEUE
      OQLT              OUT-Q LENGTH BY LTERM    (IDENTIFIER)  A   QUEUE
      OQLN              OUT-Q LENGTH BY LINE     (IDENTIFIER)  A   QUEUE
      OQND              OUT-Q LENGTH BY NODE     (IDENTIFIER)  A   QUEUE
      OUTLT             MSGS OUTPUT BY LTERM     (IDENTIFIER)  A   QUEUE
      OUTLN             MSGS OUTPUT BY LINE      (IDENTIFIER)  A   QUEUE
      OUTND             MSGS OUTPUT BY NODE      (IDENTIFIER)  A   QUEUE
      OQLK              OUTPUT QUEUE BY LINK     (LOGICAL-LINK A   QUEUE
```

Figure 10. List Data Collection Monitors Application

This application allows service selection by line command and shows how many monitors are already active, the allowable parameters for each service, the service security classification, the area of IMS being monitored, and the service status by:

Field Name	Description
LC	A line command input field. One-character line commands can be entered in this field. The line commands can be used to access a data entry panel to define the options to activate a new SET timer request, display HELP information, or show a list of active timer requests for the selected monitor (see “Line Commands” on page 90). The Active Timer Requests list can be used with line commands to view, modify, or replicate the data collection options for a selected request, as described in Chapter 6, “Displaying a List of Active Timer Requests (Primary Menu Option 2)” on page 65.
SERV	A scrollable list of all the data collection monitor services by service select code.

# ACTIVE	The number of timer requests already active for the service.
TITLE	The service title.
PARM TYPE	A short description of the parameters that can be used, if the service allows parameters.
SEC	The security code for user access to the service.
AREA	The IMS resource area being analyzed. This field could contain:

Field Data	Description
MFS	Terminal I/O
QUEUE	IMS queuing
SCHED	Scheduling of application programs in the dependent region
REGN	Application program activity in the dependent regions
DB	Database activity and buffer pool utilization
INTNL	IMS internal functions
IMVS	IMS and OS/390 interactions
LOCK	IRLM functions
IWKLD	IMS workload
IWTRN	IMS workload transactions
IWDB2	IMS workload DB2 activity
IWFP	IMS workload Fast Path activity
IWGBL	IMS workload global IMS region calls

STAT	The service status (LOCK or blank).

SORT Primary Command

When the list of data collection monitor services is displayed initially, the list is sorted by resource area. SORT can be entered in the COMMAND field of the display to sort the list by any of the following column headings. The first two characters of the column heading are used with SORT as follows:

```
SORT cc
```

where `cc` can be any of the following two characters:

SE Sorts the list alphabetically by service name (SERV column).
AC Sorts the list in a numerically descending order (# ACTIVE column).
TI Sorts the list alphabetically by service title (TITLE column).
SC Sorts the list alphabetically by the security code.
AR Sorts the list alphabetically by the resource area (AREA column) and by the service name within the area (default).
ST Sorts the list alphabetically by the service status displayed (STATUS column).

SORT without parameters sorts the list by resource area.

AREA Primary Command

You can use the AREA command to list only the services related to a specific area. The possible areas that can be specified are shown in the AREA column. For example, to list only the IMS database services, type in the COMMAND field:

```
AREA DB
```

Type AREA without an area name to return to the list of all the services.

Line Commands

Entering one of the following one-character line commands in the LC field for a service executes the line command function. Multiple selections can be entered at one time by selecting a series of services and pressing the ENTER key. Each data-collection timer request in a series is submitted by pressing the ENTER key and then the END key to process the next request.

Line Command	Description
S	SETUP. Displays a data entry panel showing the valid request options to start timer-driven data collection (see “Start Data Collection Timer Request (S Line Command)” on page 91). Each timer-driven request must be unique and is defined by the service select code and an optional parameter.
D	DISPLAY. Displays only those timer-driven requests that are active for the selected monitor, as shown by “Active Timer Requests for a Selected Monitor (D Line Command)” on page 92.
H	HELP. Displays the HELP information for the service. This shows the service title, describes the measured value, defines any parameters, and shows the format of the monitor warning message.

The following commands are for system programmer use and are restricted by a security access code:

L	LOCK. Locks this service. The service cannot be used again until it is unlocked.
U	UNLOCK. Unlocks this service. A service can be locked by the use of the LOCK command or a service ABEND.

Start Data Collection Timer Request (S Line Command)

The S line command can be used to start a data collection timer request. There are three types of data collection timer services:

- Monitor Workload Wait service
- Workload Trace service
- Monitor services

By default, the Data Collection Monitors list is sorted by AREA (see “SORT Primary Command” on page 89).

The Workload Monitors have a unique, one-character prefix (@, #, $).

Each service type has its own set of request options. When S is used to start a monitor request, a data entry panel containing the request options is displayed. The service field is preset with the code of the selected service. The input fields are prefixed with a highlighted ===> symbol. Any default values for a field are displayed.

Each request must be unique and is defined by the service select code and an optional parameter (reqid). The parameter is required if the same service is requested more than once.

The request is submitted when the ENTER key is pressed. A short message in the upper-right corner of the display shows the result of the request. If an ERROR IN REQUEST message is displayed, a short explanatory message is displayed on the third line. Pressing the END key (PF3/15) redisplays the Data Collection Monitors list.

Active Timer Requests for a Selected Monitor (D Line Command)

This list application shows each active timer request for only the service selected from the Data Collection Monitors list with the D line command. Its use is described in Chapter 6, "Displaying a List of Active Timer Requests (Primary Menu Option 2)" on page 65.

```
BMC SOFTWARE  -------------- ACTIVE TIMER REQUESTS ---------- FUNCTION COMPLETE
  COMMAND ===>                                                TGT ===> IMSA
                                    INPUT    INTVL ==> 3      TIME --  07:46:04
  COMMANDS: SM (START MONITORS), SORT, AREA, X ON|OFF, DM (DMON), DW (DWARN)
  LC CMDS:  S (SELECT), W (SHOW),      M (MODIFY),
            P (PURGE),  R (REPLICATE), H (HELP),   Z (STOP),                  >>>
  LC   SERV  PARM     TITLE                  CURRENT  WVAL  -8-6-4-2-0+2+4+6+8+
       DBIO           DB I/O COUNT BY SUBPOOL     32   30|  *********W          |
       DBIO  1        DB I/O COUNT BY SUBPOOL      0   NZ|           W          |
       DBIO  3        DB I/O COUNT BY SUBPOOL     10   10|  *********W          |
       DBIO  2        DB I/O COUNT BY SUBPOOL      0   NZ|           W          |
       **************************** END OF REQUESTS *****************************
```

Figure 11. Selected Monitor Active Timer Requests List

Chapter 9. Requesting a Monitor

These services measure and collect data about IMS resource and workload performance over time and detect when conditions exceed a user-defined threshold. You activate a service by issuing a request for it.

You can issue requests to:

- Access the monitors easily through ISPF-like menus and scrollable lists
- View plot or graphic monitor summary displays that can be refreshed in a user-defined cycle
- Start or stop a monitor from a MAINVIEW AutoOPERATOR EXEC
- Start a series of monitors automatically when the system starts or at your request
- Print a screen image of a monitor plot or graphic summary display to the online BBI-SS PAS Image log automatically, or to the TS Image log, or your BBISPRNT data set

All Workload Monitor services, except the global region call monitors, measure data from the transaction records passed to the BBI-SS PAS from the Event Collector at transaction completion (SYNC point). These records contain an identification of the specific transaction, timestamps, elapsed timings, and activity indicators. The many identification fields allow flexible workload selections by transaction, program, PSB, region, class, LTERM, user ID, program type, and transaction type.

Note: If the SAPEXIT is installed as described in the *MAINVIEW for IMS Online – Customization Guide*, the SAP identifier replaces the IMS transaction code in the transaction records. That allows the monitors to collect data based on the SAP identifier instead of the transaction code.

The data collected by these monitors is affected by the following:

- Data items from the transaction records are posted only at transaction completion (sync point). All transactions completing within a requested monitor sampling interval are included in the value calculated for that interval. This means, for example, that a program currently in a loop does ***not*** show up immediately in an increased measured response time (the transaction does not complete; therefore, it cannot be posted).
- To provide valid averages per transaction, monitors collecting Fast Path activity collect data only from transactions that made at least one Fast Path call. Monitors collecting DB2 activity collect data only from transactions that made at least one DB2 call.
- For MVIMS to calculate meaningful values over time, it can evaluate only response-type transactions, which are those running in MPP (message processing program) regions, in JMP (Java message processing program) regions, in MDP (message-driven program) regions, or under DBCTL threads. That means that BMP and JBP transactions are not measured by these monitors, because they would skew the results.

  For example, the average elapsed time of 100 MPP transactions is measured at one second. If a BMP transaction that ran for two hours (7200 seconds) were included with the MPP transactions, the average elapsed time would be 72.3 seconds. The same is true of counts, such as DB2 calls. If DB2 calls were made by this BMP transaction, these calls would be posted only when the BMP transaction completes even though they occur throughout the two-hour life of the transaction.

To avoid this skewing problem and to provide measures of total system activity, other workload monitor services measure call activity as it occurs in the regions. These services, described in Chapter 23, “Global Region Call Monitors” on page 187, are global services because they can measure MPP, JMP, BMP, JBP, and DBCTL workloads and do not require that a transaction complete before its effects can be seen. These methods are described in the following sections.

Starting a Monitor

A monitor can be started by:

- Selecting one or more services from a monitor service list
  - Access the Data Collection Monitors list with the SM application transfer command from the list of active timer requests (Primary Menu Option 2).
  - Use the S line command to select a service (see Appendix 8, “Displaying a List of Data Collection Monitors (SM Command)” on page 87).
- Replicating an active monitor request from the Active Timer list application (Primary Menu Option 2)

  You can replicate a monitor service request by using the R line command in the Active Timer list.
  - Access the Active Timer list application directly from the Primary Option Menu (Primary Menu Option 2) to view all active monitors.
  - Use the D line command from the Data Collection Monitors service list application (SM command) to access the Active Timer list for only the selected service.
- Starting a monitor service from BBPARM with other service requests

  Define a series of SET requests as a member of your BBI-SS PAS BBPARM data set that can be started automatically when the system starts or at your request (see “Multiple Requests” on page 302).
- Starting a monitor service from a MAINVIEW AutoOPERATOR EXEC

  Write an EXEC that starts a monitor service (MAINVIEW AutoOPERATOR must be installed).

  Use the IMFEXEC IMFC command followed by the service name, optional parameters, and an identifier for the target system; for example:

  ```
  IMFEXEC IMFC SET REQ=DBTOT I=00:06:00 TARGET=PROD1
  ```

Using the Resource Monitor Data Entry Panel

Each Resource Monitor is a monitoring service that measures an IMS resource over time and issues warning messages whenever user-defined thresholds are exceeded. To display the data entry panel, shown in Figure 12, for requesting a resource monitor service

1. Enter an SM command as described in Chapter 8, “Displaying a List of Data Collection Monitors (SM Command)” on page 87.
2. Use an S line command to select any of the Resource Monitor services shown in the display list of monitor services.

```
BMC SOFTWARE  ---------- START RESOURCE MONITOR REQUEST -----  PERFORMANCE MGMT
COMMAND ===>                                                  TGT ===> IMSA

                      LKREQ - IRLM LOCK REQUESTS

 PARM     ==>                             (Resource Selection Parameter)

 INTERVAL ==> 00:01:00 START ==>          STOP ==>                QIS ==> YES

 WVAL     ==>          WMSG  ==>          WLIM ==> 10  WIF ==> 1   WIN ==> 1

 RST      ==> HOT                         (Restart Option: HOT,COLD,PUR,QIS)

 TITLE    ==>                             (Title)

 PLOTMAX  ==>                             (Maximum PLOT X-Axis Value)

 RANGES   ==>                             (1-4 Range Distr. Upper Limits)

 LOG      ==>                             (NO,ATSTOP,ATPD,ATINTVL,ATWARN)
```

Figure 12. Start Resource Monitor Request Panel

These options are SET keywords for a monitor service request. These keywords are described in Table 3 on page 96 in alphabetical order.

Table 3. SET Keywords to Activate a Resource Monitor Service

Keyword	Operand	Description
INTERVAL|I	hh:mm:ss	The time interval between successive invocations of the requested service. The default is one minute (00:01:00) or as specified by the user in the BBIISP00 member of the BBPARM data set. It can be used with the LOG keyword to request automatic logging of a display to the BBI-SS PAS Image log. **Note:** The MOD keyword cannot be used to change this option. The request must be purged and a new request must be made.
LOG		Specifies if and when automatic logging occurs. Analyzer, DMON or DWARN, or PLOT displays can be logged to the BBI-SS PAS Image log. For the monitors, a PLOT of the latest data is written to the BBI-SS PAS Image log.
	NO	Specifies no logging. Default for monitor services. `LOG=NO` is the only valid option for monitors that only measure a condition against a warning threshold (data measurement type of warning only).
	ATSTOP	Causes logs to display when processing of this request is stopped. If `QIS=Y` has been specified in the request, `LOG=ATSTOP` is invoked at IMS termination and at BBI-SS PAS termination.
	ATPD	Causes logs to display at each period of 10 intervals.
	ATINTVL	Causes logs to display at each interval as specified by the user with the INTERVAL parameter or in the BBIISP00 member of the BBPARM data set. Default for analyzers. ATINTVL is the default for logging images of the analyzer services to the BBI-SS PAS Image log.
	ATWARN	Causes a plot to be logged whenever a warning message is generated by the associated monitor. NO is the default for the monitor services.
PARM	id	If you want to run multiple requests, specify an ID in the PARM field to make this request unique. The id is a unique 1- to 8-character identifier). Blank can also be used as an identifier for one request. The identifier is used to specify which monitor data is to be displayed with the plot display service.
PLOTMAX	n	Specifies the maximum value for the X-axis of a PLOT graph. Minimum is 50. The specified value is adjusted to the nearest multiple of 50. Percentages displayed by some services are always set at 100.

Table 3. SET Keywords to Activate a Resource Monitor Service (continued)

Keyword	Operand	Description
QIS		Defines the action to be taken for the service when IMS is not active.
	YES	Specifies that the service is to be quiesced. This is the default.
	NO	Specifies that the service is to start or continue running. **Note:** When `QIS=NO` is specified, monitors that require IMS continue to be scheduled at each interval; however, they return zero values. Any analyzer services set up to run asynchronously fail with a short message of CANNOT LOCATE IMS SPECIFIED in the first line. The BBI-SS PAS Image log contains screen images of these services.
RANGES	n [,n,n,n]	Up to four upper-limit values can be specified for the distribution range of any data collection monitor service. An implied limit of the maximum data measurement value is always defined internally. This information is used to produce a frequency distribution of the data measurement value at the bottom of the plot display (see "Monitor History Display (PLOT)" on page 77). If RANGES is defined, the distribution is updated at each interval with the current measurement value. A plot of the history displays this distribution. Default is no ranges. **Note:** The MOD keyword cannot be used to change this option. The request must be purged and a new request must be made.
RST		Defines the restart option to be used when a service is quiesced because of an inactive IMS subsystem or `RRR` is specified for BLK (see "Multiple Requests" on page 302). Default is HOT.
	HOT	Restarts the service automatically without loss of previous data.
	COLD	Restarts the service automatically; all previously collected data is deleted.
	PUR	Purges the service automatically when the target IMS starts.
	QIS	Keeps the service in a quiesced state until it is purged by an authorized user.

Table 3. SET Keywords to Activate a Resource Monitor Service (continued)

Keyword	Operand	Description
START	hh:mm:ss	Requests monitor start time. If the time entered is more than 10 minutes prior to the current time, 24 hours are added to the specified time and the request is started the next day. To start a request at midnight, specify 24:00:00. The default is the next full minute. **Note:** This option cannot be modified. The request must be purged and a new request must be made.
STOP	`nn|hh:mm:ss`	Requests monitor stop limit, where `nn` is length of time in minutes and `hh:mm:ss:` is a timestamp. Processing stops at the end of the last interval before the specified stop time. This time is displayed in the STOP field when the request is viewed with the R, P, M, or W line commands from the Active Timer Requests list (Primary Menu Option 2) application. If the time entered is the same as the START time, 24 hours are added to the STOP time.
TITLE|T	'c...c'	Defines a service display title and the contents of a warning message (1 to 24 characters). This user-defined title replaces the default service title. If the title is specified with SET as a single request or in a BBPARM member as a series of requests, it must be enclosed in single quotes. **Note:** The MOD keyword cannot be used to change this option. The request must be purged and a new request must be made.
WIF	n	Defines the number of times the exception is to be detected before the first message is sent. Default is 1.
WIN	n	Defines the number of times the exception is to be detected between messages. Default is 1.
WLIM	n	Defines the maximum number of warning messages to be sent for one continuous occurrence of the warning condition. Resets automatically when condition no longer exists. Default is 10.

Table 3. SET Keywords to Activate a Resource Monitor Service (continued)

Keyword	Operand	Description
WMAX|WVAL	NNE	Specifies a warning threshold. The warning condition exists if the current data measurement exceeds the defined threshold. If <n is specified, a warning is issued when the sampled value is less than or equal to the threshold.
	NZ	NZ (not zero) issues a warning when the measured value is greater than zero. Default is 0 or no warnings. **Note:** If the service measures time, the measurement units are specified in seconds and tenths of seconds (optional).
WMSG	WTO|MTO| LOG	Directs warning messages to the system console, the IMS MTO, or both. Warning messages are always written to the active BBI-SS PAS Journal log. Default is to write only to the log. An existing WTO or MTO warning message can be reset to the Journal log by using the LOG operand on a modify request. Routing and descriptor codes can be specified for WTO messages in the BBIISP00 member of the BBPARM data set.

Using the Workload Monitor Data Entry Panel

Each Workload Monitor is a monitoring service that measures a key IMS workload indicator over time and issues warning messages whenever user-defined thresholds are exceeded. To display the data entry panel, shown in Figure 13, for requesting a workload monitor service:

1. Enter an SM command as described in Chapter 8, “Displaying a List of Data Collection Monitors (SM Command)” on page 87.

2. Use an S line command to select any of the services shown in the display list of monitor services that start with a # or @ character. An S line command for a WM service displays the request data entry panel shown in Figure 13 on page 100.

```
BMC SOFTWARE ----------- START WORKLOAD MONITOR REQUEST -----  PERFORMANCE MGMT
COMMAND ===>                                                TGT ===> IMST

                       @ELAP - AVG ELAPSED TIME

PARM     ==>                             (Workload Monitor Identifier)
INTERVAL ==> 00:01:00 START ==>          STOP ==>                 QIS ==> YES
WVAL     ==>          WMSG  ==>          WLIM ==> 10  WIF ==> 1   WIN ==> 1
TITLE    ==>                             (Title)
RST      ==> HOT                         (Restart Option: HOT,COLD,PUR,QIS)
PLOTMAX  ==>                             (Maximum PLOT X-Axis Value)
RANGES   ==>                             (1-4 Range Distr. Upper Limits)
LOG      ==>                             (NO,ATSTOP,ATPD,ATINTVL,ATWARN)
Specify Workload Selections:
   TRANTYPE ==>                          (ALL, DLI, DB2, FP)
   TRAN     ==>
   CLASS    ==>
   PROG     ==>
   PSB      ==>
   REGION   ==>
   RGNID    ==>
   TERM     ==>                          (name, SYNCLOCK, MSCCLOCK)
   USERID   ==>
   PGMTYPE  ==>                          (MPP,MDP,IFP,FPU,TPI,BMP,DBT,NOTDBT)
```

Figure 13. Start Workload Monitor Request Panel

These options are SET keywords for a monitor service request. These keywords are described in Table 4 on page 101 in alphabetical order.

Table 4. SET Keywords to Activate a Workload Monitor Service

Keyword	Operand	Description
INTERVAL|I	hh:mm:ss	The time interval between successive invocations of the requested service. The default is one minute (00:01:00) or as specified by the user in the BBIISP00 member of the BBPARM data set. It can be used with the LOG keyword to request automatic logging of a display to the BBI-SS PAS Image log. **Note:** The MOD keyword cannot be used to change this option. The request must be purged and a new request must be made.
LOG		Specifies if and when automatic logging occurs. Analyzer, DMON or DWARN, or PLOT displays can be logged to the BBI-SS PAS Image log. For the monitors, a PLOT of the latest data is written to the BBI-SS PAS Image log.
	NO	No logging. Default for monitor services. **Note:** `LOG=NO` is the only valid option for monitors that only measure a condition against a warning threshold (data measurement type of warning only).
	ATSTOP	Logs display when processing of this request is stopped. If `QIS=Y` has been specified in the request, `LOG=ATSTOP` is invoked at IMS termination and at BBI-SS PAS termination.
	ATPD	Logs display at each period of 10 intervals.
	ATINTVL	Logs display at each interval as specified by the user with the INTERVAL parameter or in the BBIISP00 member of the BBPARM data set. Default for analyzers. ATINTVL is the default for logging images of the analyzer services to the BBI-SS PAS Image log
	ATWARN	Logs a plot whenever a warning message is generated by the associated monitor. NO is the default for the monitor services.
PARM	`id`	Where `id` is a unique 1- to 8-character identifier. If you want to run multiple requests, specify an ID in the PARM field to make this request unique. Blank can also be used as an identifier for one request. The identifier is used to specify which monitor data is to be displayed with the plot display service.
PLOTMAX	n	Specifies the maximum value for the X-axis of a PLOT graph. Minimum is 50. The specified value is adjusted to the nearest multiple of 50. Percentages displayed by some services are always set at 100.

Table 4. SET Keywords to Activate a Workload Monitor Service (continued)

Keyword	Operand	Description
QIS		Defines the action to be taken for the service when IMS is not active.
	YES	Specifies that the service is to be quiesced. This is the default.
	NO	Specifies that the service is to start or continue running. **Note:** When `QIS=NO` is specified, monitors that require IMS continue to be scheduled at each interval; however, they return zero values. Any analyzer services set up to run asynchronously fail with a short message of CANNOT LOCATE IMS SPECIFIED in the first line. The BBI-SS PAS Image log contains screen images of these services.
RANGES	n [,n,n,n]	Up to four upper-limit values can be specified for the distribution range of any data collection monitor service. An implied limit of the maximum data measurement value is always defined internally. This information is used to produce a frequency distribution of the data measurement value at the bottom of the plot display (see “Monitor History Display (PLOT)” on page 77). If RANGES is defined, the distribution is updated at each interval with the current measurement value. A plot of the history displays this distribution. Default is no ranges. **Note:** The MOD keyword cannot be used to change this option. The request must be purged and a new request must be made.
RST		Defines the restart option to be used when a service is quiesced because of an inactive IMS subsystem or `RRR` is specified for BLK (see “Multiple Requests” on page 302). Default is HOT.
	HOT	Restarts the service automatically without loss of previous data.
	COLD	Restarts the service automatically; all previously collected data is deleted.
	PUR	Purges the service automatically when the target IMS starts.
	QIS	Keeps the service in a quiesced state until it is purged by an authorized user.
START	hh:mm:ss	Requests monitor start time. If the time entered is more than 10 minutes prior to the current time, 24 hours are added to the specified time and the request is started the next day. To start a request at midnight, specify 24:00:00. The default is the next full minute. **Note:** This option cannot be modified. The request must be purged and a new request must be made.

Table 4. SET Keywords to Activate a Workload Monitor Service (continued)

Keyword	**Operand**	**Description**
STOP	`nn|hh:mm:ss`	Requests monitor stop limit, where `nn` is length of time in minutes and `hh:mm:ss:` is a timestamp. Processing stops at the end of the last interval before the specified stop time. This time is displayed in the STOP field when the request is viewed with the R, P, M, or W line commands from the Active Timer Requests list (Primary Menu Option 2) application. If the time entered is the same as the START time, 24 hours are added to the STOP time.
TITLE|T	'c...c'	Defines a service display title and the contents of a warning message (1 to 24 characters). This user-defined title replaces the default service title. If the title is specified with SET as a single request or in a BBPARM member as a series of requests, it must be enclosed in single quotes. **Note:** The MOD keyword cannot be used to change this option. The request must be purged and a new request must be made.
WIF	n	Defines the number of times the exception is to be detected before the first message is sent. Default is 1.
WIN	n	Defines the number of times the exception is to be detected between messages. Default is 1.
WLIM	n	Defines the maximum number of warning messages to be sent for one continuous occurrence of the warning condition. Resets automatically when condition no longer exists. Default is 10.
WMAX|WVAL	n|<n	Specifies a warning threshold. The warning condition exists if the current data measurement exceeds the defined threshold. If <n is specified, a warning is issued when the sampled value is less than or equal to the threshold.
	NZ	NZ (not zero) issues a warning when the measured value is greater than zero. Default is 0 or no warnings. **Note:** If the service measures time, the measurement units are specified in seconds and tenths of seconds (optional).

Table 4. SET Keywords to Activate a Workload Monitor Service (continued)

Keyword	Operand	Description
WMSG	WTO|MTO| LOG	Directs warning messages to the system console (WTO), the IMS MTO, or both. Warning messages are always written to the active BBI-SS PAS Journal log. Default is to write only to the log. An existing WTO or MTO warning message can be reset to the Journal log by using the LOG operand on a modify request. Routing and descriptor codes can be specified for WTO messages in the BBIISP00 member of the BBPARM data set.
Note: The following workload selection keywords act as filters to limit data collection to a specific part of the IMS workload		
CLASS	class	Qualifies a monitor request by the numeric scheduling class(es). Multiple classes can be specified as follows: `CLASS ==> 1 2 3` Default is to include all IMS scheduling classes. **Note:** This workload selection keyword cannot be used for a request for DBCTL threads.
PGMTYPE| PTYP	MPP|MDP IFP|FPU TPI|BMP DBT|NOTDBT	Qualifies a monitor request by program type. Multiple program types can be specified. The types are: MPP Message processing program or Java message processing program (JMP) MDP Message-driven program IFP IMS Fast Path program FPU Fast Path utility TPI CPI-C driven program BMP Batch message processing program or Java batch message processing program (JBP) DBT DBCTL thread (CICS or ODBA) Default is all types. DBT selects only programs that use a DBCTL thread. NOTDBT selects programs that do not use a DBCTL thread.
PROG	program name	Qualifies a monitor request by the specified program. Valid entries are 1- to 8-character alphanumeric program names, which can be used with the plus qualifier (+). Multiple program names can be specified; for example: `PROG ==> abc cde+`
PSB	psb name	Qualifies a monitor request by the specified IMS program specification block. Valid entries are 1- to 8-character alphanumeric PSB names, which can be used with the plus qualifier (+). Multiple PSB names can be specified; for example: `PSB ==> ab1 c+2`

Table 4. SET Keywords to Activate a Workload Monitor Service (continued)

Keyword	Operand	Description
REGION	job name	Qualifies a monitor request by the specified job name of an IMS region. Valid entries are 1- to 8-character alphanumeric region job names, which can be used with the plus qualifier (+). Multiple region job names can be specified; for example: `REGION ==> +ba ace`
RGNID	region number	Qualifies a monitor request by the specified region ID. Valid entries are 1- to 3-character numeric region IDs from 1 to `nnn`, where `nnn` is a valid region number. Multiple region IDs can be specified; for example: `RGNID ==> 1 23 200`
TERM	lterm	Qualifies a monitor request by the name of the IMS LTERM. Valid entries are 1- to 8-character alphanumeric LTERM names, which can be used with the plus qualifier (+). Multiple LTERM names can be specified; for example: `TERM ==> L071 L0+`
	SYNCLOCK	Defines a monitor request to include *all* transactions from systems synchronized to the clock of the local IMS (in other words, all local transactions, all local MSC transactions, and all nonlocal MSC transactions that arrive from a system with its clock synchronized to the clock of the local IMS).
	MSCCLOCK	Defines a monitor request to include *only nonlocal* transactions from MSC systems synchronized to the clock of the local IMS.
TRAN	trancode	Qualifies a monitor request by the specified IMS or CICS transaction. Valid entries are 1- to 8-character alphanumeric transaction codes, which can be used with the plus qualifier (+). Multiple transaction codes can be specified; for example: `TRAN ==> TR+ PAY1` For a CICS DBCTL transaction, the CICS transaction code is used. For a Fast Path transaction, the routing code assigned by the DBFHAGU0 exit routine is used.

Table 4. SET Keywords to Activate a Workload Monitor Service (continued)

Keyword	**Operand**	**Description**
TRANTYPE| TTYP	ALL|DB2| DLI|FP	Qualifies a monitor request by transaction type: ALL All transaction types DB2 Those that access DB2 DLI Those that make DL/I database calls FP Those that make Fast Path database calls The default is all types, and multiple transaction types can be specified.
USERID	userid	Qualifies a monitor request by the name of the IMS sign-on user ID. Valid entries are 1- to 8-character alphanumeric IMS sign-on user IDs, which can be used with the plus qualifier (+). Multiple user IDs can be specified; for example: `USERID ==> USER1 USER2 USER6`

Using the Workload Global Region Call Monitor Data Entry Panel

Each Workload Global Region Call Monitor is a monitoring service that measures region, database, and message queue calls over time and issues warning messages whenever user-defined thresholds are exceeded. To display the data entry panel, shown in Figure 14, for requesting a workload global region call monitor service:

1. Enter an SM command as described in Chapter 8, “Displaying a List of Data Collection Monitors (SM Command)” on page 87.
2. Use an S line command to select any of the services shown in the display list of monitor services that start with a $ character. An S line command for a global region call monitor service displays the request data entry panel shown in Figure 14.

```
BMC SOFTWARE  ---------- START WM GLOBAL MONITOR REQUEST ----- PERFORMANCE MGMT
COMMAND ===>                                              TGT ===> IMSA

                      $CTOT - ALL REGION CALLS

 PARM     ==>                            (Workload Monitor Identifier)

 INTERVAL ==> 00:01:00 START ==>          STOP ==>                QIS ==> YES

 WVAL     ==>          WMSG  ==>          WLIM ==> 10  WIF ==> 1   WIN ==> 1

 TITLE    ==>                            (Title)

 RST      ==> HOT                        (Restart Option: HOT,COLD,PUR,QIS)

 PLOTMAX  ==>                            (Maximum PLOT X-Axis Value)

 RANGES   ==>                            (1-4 Range Distr. Upper Limits)

 LOG      ==>                            (NO,ATSTOP,ATPD,ATINTVL,ATWARN)

 Specify Workload Selections:
   CALLTYPE ==>                          (ALL, MSG, DB2, DLI)
```

Figure 14. Start Workload Global Region Call Monitor Request Panel

These services use the same options as the workload monitor services except for the workload selection criteria. The CALLTYPE option allows workload selection by type of calls for the region call monitor services as described in Table 5.

Table 5. SET Keyword to Define Workload Selection for Global Region Call Monitors

Keyword	**Operand**	**Description**
CALLTYPE| CTYP	<u>ALL</u> MSG DB2 DLI	Specifies the type of calls to be collected by $CMPP, $CBMP, $CTOT, and $CDBT. You can specify the following: ALL All calls MSG Message queue calls DB2 DB2 calls DLI DL/I calls

Displaying Monitor Data

A display of the data collected by monitors can be requested by:

- Selecting a scrollable list of active monitors and their current values:
  - Access the Active Timer Request list application directly from the Primary Option Menu (Primary Menu Option 2) to view all active monitors.
  - Use the D line command from the Data Collection Monitors service list application (SM command) to access the Active Timer list for only the selected service.
  - Selecting an active monitor summary display (DMON)
  - Use the DM application transfer command.
  - Invoke the DMON service from the SERV field of any display with a SET request.
  - Selecting an active monitor warning display (DWARN)
  - Use the DW application transfer command.
  - Invoke the DWARN service from the SERV field of any display with a SET request.
  - Selecting a graphic plot of the historical data collected by one monitor
  - Use the S line command in the Active Timer list (Primary Menu Option 2) to select a plot of the data collected by that active monitor.
  - Move the cursor to one of the monitor requests in a DMON or DWARN display and press ENTER to view a plot of data collected by that monitor as described in “Active Monitor Summary Display (DMON)” on page 83 and “Active Monitor Warning Display (DWARN)” on page 85.
  - Invoke the PLOT service from the SERV field of any display with a SET request.
- Setting up monitor graphic displays for timed, cyclic refresh

  Select Option C, CYCLE SETUP, from the Primary Option Menu to set up a continuous timed cycle of refreshable monitor plot (PLOT, with a service name parameter) or graphic summary displays (DMON or DWARN). For more information about this option, see the *Using MAINVIEW* manual.

Logging Monitor Data

An image of a monitor plot or graphic summary display (DMON or DWARN) can be recorded in the BBI-SS PAS Image log automatically, to your TS Image log, or to a BBISPRNT data set at your request by:

- Logging an image to the BBI-SS Image log
  - Specify the LOG option when starting the monitor or modifying it.
  - Write an EXEC that defines a SET log request for a monitor plot or graphic summary display (DMON or DWARN) (MAINVIEW AutoOPERATOR must be installed).

Use the IMFEXEC IMFC command followed by the service name, an optional parameter, and an identifier for the target system as in the following examples:

```
IMFEXEC IMFC SET REQ=PIENQ WMAX=80,LOG=ATWARN TARGET=TEST1
```

```
IMFEXEC IMFC DMON PIENQ I=00:05:00 TARGET=TEST1
```

  - Define a SET log request for a graphic summary display (DMON or DWARN) as a member of your BBI-SS PAS BBPARM data set that can be started automatically when the system starts or at your request (look up "grouping requests" in the *Analyzers Reference Manual* index).

    BBSAMP member ILOGJCL can be used to create a hardcopy of your Image log data sets.

- Logging a display image record to the TS Image log

  Enter a Y in the LOG field of the plot or graphic summary display (DMON or DWARN) to record the image in your TS Image log.

- Logging a screen image to your BBISPRNT data set

  Press the PF4/16 key to record a plot or graphic summary display image in your BBISPRNT data set.

  BBSAMP member SLOGJCL can be used to create a hardcopy of your BBISPRNT data set.

Stopping a Monitor

A monitor service or Image log request can be stopped by:

- Stopping the monitor request with a Z line command

  Use the Z line command from the Active Timer list application as described in Chapter 6, "Displaying a List of Active Timer Requests (Primary Menu Option 2)" on page 65.

- Setting a stop time for automatic completion of data collection

  Specify the STOP value (as a timestamp or interval count) on the Start or Modify panel for the monitor. The collected data remains available for viewing until the monitor is purged.

- Purging a request from the Active Timer list with a P line command

  Use the P line command from the Active Timer list application as described in Chapter 6, "Displaying a List of Active Timer Requests (Primary Menu Option 2)" on page 65.

- Purging a service request with a SET request

  Issue a PRG request with SET from the SERV field of any display, a BBPARM member (see "grouping requests" in the *Analyzers Reference Manual* index), or a MAINVIEW AutoOPERATOR EXEC (MAINVIEW AutoOPERATOR must be installed);
  for example:

  ```
  SERV ==> SET
  PARM ==> PRG=reqid|ALL
  ```

- Stopping a service with a SET request

  Enter a STOP time in the data entry panel for the monitor request or use a SET request in the SERV field of any display as follows:

  ```
  SERV ==> SET
  PARM ==> REQ=DBTOT,START=00:11:00,STOP=00:12:00,LOG=ATSTOP,I=00:01:00
  ```

  The request starts at 11 minutes after midnight and stops 12 minutes after midnight; the plot display is logged to the SS Image log when the request stops. STOP or STOPCNT (see Table 16 on page 308) can be used with a SET request.

Qualifying Monitor Requests

Requests for multiple resources with similar names can be made by using a + character as a name qualifier. The + character can be used generically or positionally:

- As a generic resource name qualifier, it cannot be followed by any other character.
- As a positional qualifier, it must be repeated for every character to be replaced.

If you receive a message that contains the following in the BBI-SS PAS journal:

```
IM4704E ... INVALID PARM - R/C 04
```

the qualified resource could not be found.

Resource Monitors

A qualifier can be used when making Resource Monitor service requests with the following resource parameters:

Resource Parameter	Resource Monitor Service
Program names	ARVBG, ARVPR, INQBG, PRCBG, PRCPR
Transaction names	ARVTR, INQTR, PRCTR
Lterm names	OQLT, OUTLT
Node names	OQND, OUTND
Region names	PIENQ

For example, if you enter the following in the data entry panel for the PIENQ service:

```
PARM ==> ABC+
```

the PI enqueues held by a group of regions with names starting with ABC are sampled.

Entering the following data collection request in the SERV field and a qualified resource name in the PARM field of the Service Display panel:

```
SET
REQ=PIENQ A++D
```

monitors PI enqueues held by regions starting with an A character, followed by any two characters, and ending with a D.

Workload Monitors

Qualifiers can be used with Workload Monitor service requests. For example, if you enter the following in the data entry panel for the @ELAP service:

```
TRAN ==> PAY+
```

only transactions with names that start with PAY are sampled.

Entering the following data collection request in the SERV field and a qualified resource name in the PARM field of the Service Display panel:

```
SET
REQ=@ELAP A+++++C
```

monitors IMS elapsed time of all transaction codes starting with an A character, followed by any five characters, and ending with a C.

Chapter 10. Resource Monitor Services (Quick Reference)

This section provides an alphabetical list of all the Resource Monitor services and their parameters with page references to a more detailed description about their use. See Chapter 8, “Displaying a List of Data Collection Monitors (SM Command)” on page 87 for a complete description of the Data Collection Monitors list application.

Table 6. Resource Monitor Service Select Codes

Service Select Code	**Parameter**	**See**
ARVBG	[psbname]	“ARVBG - Transactions Arrivals by Balancing Group (BALG)” on page 133
ARVCL	[clsnumbr]	“ARVCL - Transaction Arrivals by Class” on page 134
ARVPR	[progname]	“ARVPR - Transaction Arrivals by Program” on page 134
ARVTR	[tranname]	“ARVTR - Transaction Arrivals by Transaction Code” on page 135
CSAFR		“CSAFR - CSA Fragmentation” on page 165
CSAUT		“CSAUT - CSA Percentage of Utilization” on page 165
D2CON	[subsysid]	“D2CON - IMS Region Connection to DB2 Subsystem” on page 139
D2SON	[subsysid]	“D2SON - DB2 Sign On by Subsystem” on page 139
D2THD	[subsysid]	“D2THD - Active IMS Region Threads to DB2 Subsystem” on page 140
DBGU	[jobname]	“DBGU - Database Calls per Message Get Unique by Region” on page 140
DBHIT	[subpool number]	“DBHIT - Hit Ratio for OSAM Buffer Pool” on page 146
DBIO	[n]	“DBIO - Database I/O Count by Subpool” on page 145
DBSTL	[n]	“DBSTL - Database Buffer Steals by Subpool” on page 147
DBTOT	[jobname]	“DBTOT - Database Calls per Scheduling by Region” on page 140
DBWP		“DBWP - Database Work Area Pool Percentage of Utilization” on page 153
DEADQ	[n]	“DEADQ - Dead Letter Queue Count” on page 128
DLIDB	[NOTBMP| jobname]	“DLIDB - DL/I DB Call Time” on page 141

Table 6. Resource Monitor Service Select Codes (continued)

Service Select Code	Parameter	See
DLIDC	[GU|GUNOTWFI|NOTWFI|jobname]	"DLIDC - DL/I DC Call Time" on page 142
DLIO	[ddname]	"DLIO - DL/I EXCP Count by ddname" on page 166
DMBP		"DMBP - DMB Pool Percentage of Utilization" on page 153
DPAGE	[jobname]	"DPAGE - Demand Paging by Region" on page 166
DSAP		"DSAP - Dynamic SAP Percentage of Utilization" on page 154
ECSAU		"ECSAU - Extended CSA Percent Utilization" on page 167
EPCB		"EPCB - EPCB Pool Percentage of Utilization" on page 154
HPACC	[subpool number]	"HPACC - Hiperspace Access by Subpool" on page 147
HPHIT	[subpool number]	"HPHIT - Hiperspace Hit Ratio by Subpool" on page 148
HPSTL	[subpool number]	"HPSTL - Hiperspace Buffer Steals by Subpool" on page 149
INLK	[logical link number]	"INLK - Input Messages by Link" on page 126
INQBG	[progname]	"INQBG - Input Queue Length by Balancing Group (BALG)" on page 122
INQCL	[clsnumbr]	"INQCL - Input Queue Length by Class" on page 122
INQTR	[tranname]	"INQTR - Input Queue Length by Transaction Code" on page 123
IQSCL	[clsnumbr]	"IQSCL - Schedulable Input Queue by Class" on page 123
LAWT	[type]	"LAWT - Average Latch Wait Time" on page 155
LDLCK		"LDLCK - Number of Deadlocks" on page 169
LGMSG		"LGMSG - Long Message Queue Percentage Utilization" on page 124
LHELD	[IMS online region number]	"LHELD - Number of Locks Held" on page 169

Table 6. Resource Monitor Service Select Codes (continued)

Service Select Code	Parameter	See
LKMAX	[MPP|BMP|IFP|DBT]	"LKMAX - Maximum Locks Held by Region" on page 170
LKREQ	[GLOBAL]	"LKREQ - Number of Lock Requests" on page 171
LMAWT		"LMAWT - Maximum Average Latch Wait Time" on page 156
LSUSP		"LSUSP - Number of Suspensions" on page 171
LWAIT	[MPP|BMP|IFP|DBT]	"LWAIT - Region in IRLM Suspend" on page 172
LWNUM	[GBL|GBLssss|ssss]	"LWNUM - Number of Suspended IRLM Requests" on page 173
MFSFD		"MFSFD - Percentage of MFS Blocks Found in Pool" on page 119
MFSIO		"MFSIO - MFS I/O" on page 119
MFSIR		"MFSIR - MFS Immediate Requests" on page 120
MFSP		"MFSP - MFS Pool Percentage of Utilization" on page 156
MSGGU	[jobname]	"MSGGU - Message Calls per Message Get Unique by Region" on page 143
MSGT	[jobname]	"MSGT - Message Calls per Scheduling by Region" on page 143
OBUFW		"OBUFW - OLDS Buffer Waits" on page 157
OCHKW		"OCHKW - OLDS Check Writes" on page 157
OQLK	[logical link number]	"OQLK - Output Queue by Link" on page 126
OQLN	[lnnumber]	"OQLN - Output Queue Length by Line" on page 129
OQLT	[ltrmname]	"OQLT - Output Queue Length by LTERM" on page 129
OQND	[nodename]	"OQND - Output Queue Length by Node" on page 130
OUTLK	[logical link number]	"OUTLK - Output Messages by Link" on page 127
OUTLN	[lnnumber]	"OUTLN - Messages Output by Line" on page 130

Table 6. Resource Monitor Service Select Codes (continued)

Service Select Code	Parameter	See
OUTLT	[ltrmname]	"OUTLT - Messages Output by LTERM" on page 131
OUTND	[nodename]	"OUTND - Messages Output by Node" on page 131
PAGE	[jobname]	"PAGE - Paging (Region)" on page 167
PIENQ	[jobname]	"PIENQ - Program Isolation Enqueues by Region" on page 158
PIMAX	[MPP|BMP|IFP|DBT]	"PIMAX - Maximum Program Isolation Enqueues by Region" on page 159
PIPL		"PIPL - Program Isolation Pool Percentage of Utilization" on page 159
POOLA	[id|CSA|ECSA]	"POOLA - Pool Allocated Storage" on page 160
POOLN	[id]	"POOLN - Net Expansion Count" on page 161
POOLT	[id]	"POOLT - Total Expansion/Compression Count" on page 162
PRCBG	[psbname]	"PRCBG - Transactions Processed by Balancing Group (BALG)" on page 135
PRCCL	[clsnumbr]	"PRCCL - Transactions Processed By Class" on page 136
PRCPR	[psbname]	"PRCPR - Transactions Processed by Program" on page 136
PRCTR	[tranname]	"PRCTR - Transactions Processed by Transaction Code" on page 137
PSBP		"PSBP - PSB Pool Percentage of Utilization" on page 163
PSBW		"PSBW - PSB Work Area Pool Percentage of Utilization" on page 163
PTBLK	[GLOBAL]	"PTBLK - Number of PTB Locks" on page 173
QBLKS		"QBLKS - Queue Blocks Percentage of Utilization" on page 124
QIO		"QIO - Queue I/O" on page 124
QWAIT		"QWAIT - Queuing Waits" on page 125
RECA		"RECA - RECA Pool Percentage of Utilization" on page 164
SBUSE	[region name or blank]	"SBUSE - Sequential Buffering Storage by Region" on page 150

Table 6. Resource Monitor Service Select Codes (continued)

Service Select Code	**Parameter**	**See**
SCHFL	[PGM| OTHER| INTENT| PRIORITY| NOOTHER]	"SCHFL - Scheduling Failures by Type" on page 138
SHMSG		"SHMSG - Short Message Queue Percentage Utilization" on page 125
SYSIO	[ddname]	"SYSIO - EXCP Count by ddname" on page 168
VDBIO	[n]	"VDBIO - VSAM Database I/O by Subpool" on page 150
VDBWR	[n]	"VDBWR - VSAM Writes by Subpool" on page 151
VHIT	[subpool number]	"VHIT - VSAM Hit Ratio by Subpool" on page 152
VSEND		"VSEND - Number of VTAM Sends" on page 174
WADIO		"WADIO - WADS I/O" on page 164
WAIT		"WAIT - Region in a Long PI Wait" on page 144
WKAP		"WKAP - General Work Area Pool Percentage of Utilization" on page 164

Chapter 11. IMS MFS Services

The MFS services measure MFS I/O and MFS requests. For example, to activate the MFSIO service to track the amount and rate of I/O to the MFS format data set and log a plot of the collected data to the BBI-SS PAS Image log every hour, select the MFSIO monitor from the monitor list (SM command) and specify:

```
INTERVAL ==> 00:06:00

LOG ==> ATPD
```

MFSFD - Percentage of MFS Blocks Found in Pool

Select Code:	MFSFD
Parameter:	None
Measurement:	Percent MFS immediate block requests satisfied in pool (this is a cumulative percentage based on the totals as measured at that time) `= (immediate block requests - immediate block reads) * 100 / immediate block requests`
Data Type:	PERCENT
Default Title:	% MFS BLOCKS FOUND IN POOL
Warning:	`RM0030W (nn) hh:mm:ss title(TOTAL) = v (>thrshld) *****`

MFSIO - MFS I/O

Select Code:	MFSIO
Parameter:	None
Measurement:	Total MFS I/O performed in specified time interval `= directory reads + immediate block reads`
Data Type:	COUNT
Default Title:	MFS I/O
Warning:	`RM0010W (nn) hh:mm:ss title(TOTAL) = v IN x intrvl (>thrshld) *****`

MFSIR - MFS Immediate Requests

Select Code:	MFSIR
Parameter:	None
Measurement:	Total MFS immediate requests in specified time interval
Data Type:	COUNT
Default Title:	MFS IMMEDIATE REQUESTS
Warning:	`RM0020W (nn) hh:mm:ss title(TOTAL) = v IN x intrvl (>thrshld) *****`

Chapter 12. IMS Queuing Services

These services measure the queuing of transactions for processing. For example, to activate the INQTR service to track input queue length, issue a warning message (5 maximum) to the MTO if the queue length is greater than 20 transactions, and log a plot of the collected data to the BBI-SS PAS Image log when the message is generated, select the INQTR monitor from the monitor list (SM command) and specify:

```
WMSG ==> MTO
WVAL ==> 20
WLIM ==> 5
LOG  ==> ATWARN
```

Note: Resource Monitor uses the IMS counts-per-transaction to measure transaction arrivals and transactions processed. These IMS counters have a maximum value of 32,767 and are reset to zero only at an IMS cold start or counter overflow. If an overflow occurs for a specific transaction, IMS resets the counter to zero, which can affect the accuracy of the MVIMS counts. When MVIMS detects this condition and cannot calculate a true count, an estimate is made by computing the average of the last 10 counts.

Input Queue

The following services provide performance measurement information about the input queue length, including the number of transactions on the input queue and the number of transactions available for scheduling.

INQBG - Input Queue Length by Balancing Group (BALG)

Select Code:	INQBG
Parameter:	Program name If no parameter is entered, total input queue length for all BALGs is measured. Enter program name to measure input queue length for one BALG. Enter qualified name to measure input queue length for a group of BALGs. If you receive a message that contains the following in the BBI-SS PAS journal: `IM4704E ... INVALID PARM - R/C 04` the qualified resource could not be found.
Measurement:	Number of transactions queued when sampled: total, by BALG, or by BALG group
Data Type:	STATUS
Default Title:	IN-Q LENGTH BY BALG
Warning:	`RM0060W (nn) hh:mm:ss title(parm) = v (>thrshld) *****`

INQCL - Input Queue Length by Class

Select Code:	INQCL
Parameter:	Class number If no parameter is entered, total input queue length is measured (including Fast Path transactions). Enter a class number to measure input queue length for one processing class.
Measurement:	Number of transactions queued when sampled: total or by class
Data Type:	STATUS
Default Title:	IN-Q LENGTH BY CLASS
Warning:	`RM0080W (nn) hh:mm:ss title(parm) = v (>thrshld) *****`

INQTR - Input Queue Length by Transaction Code

Select Code: INQTR

Parameter: Transaction name

If no parameter is entered, total input queue length is measured. (This includes Fast Path transactions.)

Enter transaction name to measure input queue length for one transaction. Messages queued to Fast Path with this transaction name are not counted. See “INQBG - Input Queue Length by Balancing Group (BALG)” on page 122 for Fast Path.

Enter qualified name to measure input queue length for a group of transactions. If you receive a message that contains the following in the BBI-SS PAS journal:

```
IM4704E ... INVALID PARM - R/C 04
```

the qualified resource could not be found.

Measurement: Number of transactions queued when sampled: total, by transaction, or by transaction group

Data Type: STATUS

Default Title: IN-Q LENGTH BY TRANCODE

Warning: `RM0070W (nn) hh:mm:ss title(parm) = v (>thrshld) *****`

IQSCL - Schedulable Input Queue by Class

Select Code: IQSCL

Resource Monitor does not provide a Schedulable Input Queue by Balancing Group service for Fast Path users because all messages on a BALG are always schedulable. Use INQBG to monitor BALGs.

Parameter: Class number

If no parameter is entered, total input queue length of all transactions available for scheduling is measured (including Fast Path transactions).

Enter a class number to measure this input queue length for one processing class.

Measurement: Number of queued transactions that are available for scheduling when sampled (not locked, stopped, or priority zero): total or by class

Data Type: STATUS

Default Title: SCHDULABLE IN-Q BY CLASS

Warning: `RM0090W (nn) hh:mm:ss title(parm) = v (>thrshld) *****`

Message Queue

The following services provide performance measurement information about the message queue, which includes queue I/O and waits and usage of the long and short message queue data sets and queue blocks data set.

LGMSG - Long Message Queue Percentage Utilization

Select Code:	LGMSG
Parameter:	None
Measurement:	Long message queue data set percent utilization
	`= records allocated * 100 / total number of records`
Data Type:	PERCENT
Default Title:	LONG MSG QUEUE % UTIL
Warning:	`RM0580W (nn) hh:mm:ss title(TOTAL) = v (>thrshld) *****`

QBLKS - Queue Blocks Percentage of Utilization

Select Code:	QBLKS
Parameter:	None
Measurement:	QBLKS message queue data set percent utilization
	`= records allocated * 100 / total number of records`
Data Type:	PERCENT
Default Title:	QBLKS % UTILIZATION
Warning:	`RM0590W (nn) hh:mm:ss title(TOTAL) = v (>thrshld) *****`

QIO - Queue I/O

Select Code:	QIO
Parameter:	None
Measurement:	Total queue I/O performed in specified time interval = `reads + writes`
Data Type:	COUNT
Default Title:	QUEUE I/O
Warning:	`RM0040W (nn) hh:mm:ss title(TOTAL) = v IN x intrvl (>thrshld) *****`

QWAIT - Queuing Waits

Select Code: QWAIT

Parameter: None

Measurement: Total queuing waits in specified time interval

```
= waits for an available buffer
+ waits for other DECB to read
+ waits for other DECB to write
+ waits for purge
+ waits for buffer ENQ/DEQ
```

Data Type: COUNT

Default Title: QUEUING WAITS

Warning: `RM0050W (nn) hh:mm:ss title(TOTAL) = v IN x intrvl (>thrshld) *****`

SHMSG - Short Message Queue Percentage Utilization

Select Code: SHMSG

Parameter: None

Measurement: Short message queue data set percent utilization

```
= records allocated * 100 / total number of records
```

Data Type: PERCENT

Default Title: SHORT MSG QUEUE % UTIL

Warning: `RM0570W (nn) hh:mm:ss title(TOTAL) = v (>thrshld) *****`

MSC Links

The following services provide performance measurement information about messages received and sent on MSC logical links and the MSC logical link output queue length.

INLK - Input Messages by Link

Select Code: INLK

Parameter: Logical Link number

If no parameter is entered, the number of messages received on all MSC Logical Links is measured.

Enter Logical Link number to measure the number of messages received on one link.

Measurement: Number of messages received on the logical link(s) within the specified interval

Data Type: COUNT

Default Title: INPUT MESSAGES BY LINK

Warning: `RM0790W (nn) hh:mm:ss title(parm) = v IN x intrvl (>thrshld) *****`

OQLK - Output Queue by Link

Select Code: OQLK

Parameter: Logical Link number

If no parameter is entered, output queue length for all MSC Logical Links is measured.

Enter Logical Link number to measure the output queue length for one link.

Measurement: Output queue length for MSC Logical Links at the time of the sample; calculated as the difference between the enqueue and dequeue counts for remote transactions and MSNAMES using this logical link

Data Type: STATUS

Default Title: OUTPUT QUEUE BY LINK

Warning: `RM0770W (nn) hh:mm:ss title(parm) = v  (>thrshld) *****`

OUTLK - Output Messages by Link

Select Code:	OUTLK
Parameter:	Logical Link number If no parameter is entered, the number of messages sent on MSC Logical Links is measured. Enter Logical Link number to measure the number of messages sent on one link.
Measurement:	Number of messages sent on logical link(s) within the specified interval
Data Type:	COUNT
Default Title:	OUTPUT MESSAGES BY LINK
Warning:	`RM0780W (nn) hh:mm:ss title(parm) = v IN x intrvl (>thrshld) *****`

Output Queue

The following services provide performance measurement information about the number of output messages queued by resource, the number of messages dequeued by resource, and the number of messages queued to inactive ETO user structures.

DEADQ - Dead Letter Queue Count

Select Code: DEADQ

Parameter: Number of days or blank

Measurement: The number of messages queued to inactive ETO user structures

A user structure is considered inactive if it has not been connected to a NODE for at least as many days as specified in the parameter field. If the parameter field is left blank, the value specified for the IMS dead letter queue time is used.

The parameter tests how long the user structure has been inactive, not how long the message has been on the queue. All messages queued to a user structure are counted regardless of when they were put on the queue.

Data type: STATUS

Default Title: DEAD LETTER QUEUE COUNT

Warning: `RM0990W (nn) hh:mm:ss title(parm) = v (>thrshld) *****`

Recommendations

- Use this monitor to check for an excessive number of messages that cannot be delivered to users. When this value gets too high, select the USER display service and sort by the AGE column (number of days) to find the users that have been inactive the longest.

  By using the parameter, you can check the number of undelivered messages independent of the dead letter queue timeout value chosen. This allows you to analyze the problem before the user structure goes into a DEADQ status.

- You may want to set several monitors using different parameter values so that you can see if the queue depth tends to stabilize over time.
- To remove undelivered messages from your IMS system, either assign the LTERM to a valid USER structure or use the IMS /DEQ PURGE command.
- Messages that stay on the queue for days increase short and long message queue usage. If the DEADQ monitor shows a large number of undelivered messages, verify your message queue usage with the SHMSG and LGMSG monitors.

OQLN - Output Queue Length by Line

Select Code:	OQLN
Parameter:	Line number
	If no parameter is entered, total output queue length is measured.
	Enter line number to measure output queue length for one physical line.
Measurement:	Number of output messages queued when sampled: total or by line
Data Type:	STATUS
Default Title:	OUT-Q LENGTH BY LINE
Warning:	`RM0110W (nn) hh:mm:ss title(parm) = v (>thrshld) *****`

OQLT - Output Queue Length by LTERM

Select Code:	OQLT
Parameter:	LTERM name
	If no parameter is entered, total output queue length is measured.
	Enter LTERM name to measure output queue length for one LTERM.
	Enter qualified LTERM name to measure output queue length for a group of LTERMs. If you receive a message that contains the following in the BBI-SS PAS journal:
	`IM4704E ... INVALID PARM - R/C 04`
	the qualified resource could not be found.
Measurement:	Number of output messages queued when sampled: total, by LTERM, or by LTERM group
	Note: OQLT, OQND, OQLN, OUTLT, OUTND, and OUTLN do not include messages sent to terminals by Fast Path. These message types bypass the IMS message queue.
Data Type:	STATUS
Default Title:	OUT-Q LENGTH BY LTERM
Warning:	`RM0100W (nn) hh:mm:ss title(parm) = v (>thrshld) *****`

OQND - Output Queue Length by Node

Select Code: OQND

Parameter: VTAM node name

If no parameter is entered, total output length is measured.

Enter VTAM node name to measure output queue length for one VTAM node.

Enter qualified VTAM node name to measure output queue length for a group of VTAM nodes. If you receive a message that contains the following in the BBI-SS PAS journal:

```
IM4704E ... INVALID PARM - R/C 04
```

the qualified resource could not be found.

Measurement: Number of output messages queued when sampled: total, by VTAM node, or by VTAM node group

Data Type: STATUS

Default Title: OUT-Q LENGTH BY NODE

Warning: `RM0120W (nn) hh:mm:ss title(parm) = v (>thrshld) *****`

OUTLN - Messages Output by Line

Select Code: OUTLN

Parameter: Line Number

If no parameter is entered, total messages output are measured.

Enter physical line number to measure messages output to one line.

Measurement: Messages output (dequeues) in specified interval: total or by line

Data Type: COUNT (see note in page 121)

Default Title: MSGS OUTPUT BY LINE

Warning: `RM0140W (nn) hh:mm:ss title(parm) = v IN x intrvl (>thrshld) *****`

OUTLT - Messages Output by LTERM

Select Code: OUTLT

Parameter: LTERM name

If no parameter is entered, total messages output are measured.

Enter LTERM name to measure messages output to one LTERM.

Enter qualified name to measure messages output to a group of LTERMs. If you receive a message that contains the following in the BBI-SS PAS journal:

```
IM4704E ... INVALID PARM - R/C 04
```

the qualified resource could not be found.

Measurement: Messages output (dequeues) in specified interval: total, by LTERM, or by LTERM group

Data Type: COUNT (see note on page 121)

Default Title: MSGS OUTPUT BY LTERM

Warning: `RM0130W (nn) hh:mm:ss title(parm) = v IN x intrvl (>thrshld) *****`

OUTND - Messages Output by Node

Select Code: OUTND

Parameter: VTAM node name

If no parameter is entered, total messages output are measured.

Enter VTAM node name to measure messages output to one VTAM node.

Enter qualified name to measure messages output to a group of VTAM nodes. If you receive a message that contains the following in the BBI-SS PAS journal:

```
IM4704E ... INVALID PARM - R/C 04
```

the qualified resource could not be found.

Measurement: Messages output (dequeues) in specified interval: total, by VTAM node, or by VTAM node group

Data Type: COUNT (see note on page 121)

Default Title: MSGS OUTPUT BY NODE

Warning: `RM0150W (nn) hh:mm:ss title(parm) = v IN x intrvl (>thrshld) *****`

Chapter 13. IMS Scheduling Services

These services measure transaction arrivals and transactions processed. For example, to activate the ARVTR service to track the transaction arrival rate and log a plot of the collected data to the BBI-SS PAS Image log every hour at the end of 10 six-minute intervals, select the ARVTR monitor from the monitor list (SM command) and specify:

```
INTERVAL ==> 00:06:00
LOG      ==> ATPD
```

As long as the request is active, an online PLOT is available.

Note: Resource Monitor uses the IMS counts-per-transaction to measure transaction arrivals and transactions processed. These IMS counters have a maximum value of 32,767 and are reset to zero only at an IMS cold start or counter overflow. If an overflow occurs for a specific transaction, IMS resets the counter to zero, which can affect the accuracy of the MVIMS counts. When MVIMS detects this condition and cannot calculate a true count, an estimate is made by computing the average of the last 10 counts.

ARVBG - Transactions Arrivals by Balancing Group (BALG)

Select Code: ARVBG

Parameter: Program (PSB) name

If no parameter is entered, total transaction arrivals for all BALGs are measured.

Enter program name to measure arrivals for one BALG.

Enter qualified name to measure arrivals for a group of BALGs. If you receive a message that contains the following in the BBI-SS PAS journal:

```
IM4704E ... INVALID PARM - R/C 04
```

the qualified resource could not be found.

Measurement: Transaction arrivals (enqueues) in specified interval: total, by BALG, or by BALG group

Data Type: COUNT

Default Title: TRAN ARRIVALS BY BALG

Warning: `RM0160W (nn) hh:mm:ss title(parm) = v IN x intrvl (>thrshld) *****`

ARVCL - Transaction Arrivals by Class

Select Code:	ARVCL
Parameter:	Class If no parameter is entered, total transaction arrivals are measured (including Fast Path transaction arrivals). Enter class to measure arrivals for one processing class.
Measurement:	Transaction arrivals (enqueues) in specified interval: total or by class
Data Type:	COUNT (see note on page 121)
Default Title:	TRAN ARRIVALS BY CLASS
Warning:	`RM0180W (nn) hh:mm:ss title(parm) = v IN x intrvl (>thrshld) *****`

ARVPR - Transaction Arrivals by Program

Select Code:	ARVPR
Parameter:	Program (PSB) name If no parameter is entered, total transaction arrivals are measured (including Fast Path transaction arrivals). Enter program name to measure arrivals for one program (including Fast Path transaction arrivals). Enter qualified name to measure arrivals for a group of programs (including Fast Path transaction arrivals). If you receive a message that contains the following in the BBI-SS PAS journal: `IM4704E ... INVALID PARM - R/C 04` the qualified resource could not be found.
Measurement:	Transaction arrivals (enqueues) in specified interval: total, by program, or by program group
Data Type:	COUNT (see note on page 121)
Default Title:	TRAN ARRIVALS BY PGM
Warning:	`RM0190W (nn) hh:mm:ss title(parm) = v IN x intrvl (>thrshld) *****`

ARVTR - Transaction Arrivals by Transaction Code

Select Code: ARVTR

Parameter: Transaction name

If no parameter is entered, total transaction arrivals are measured (this includes Fast Path transactions).

Enter transaction name to measure arrivals for one transaction. For Fast Path, message arrivals with this transaction name are not counted. See "ARVBG - Transactions Arrivals by Balancing Group (BALG)" on page 133 for Fast Path.

Enter qualified name to measure arrivals for a group of transactions. If you receive a message that contains the following in the BBI-SS PAS journal:

```
IM4704E ... INVALID PARM - R/C 04
```

the qualified resource could not be found.

Measurement: Transaction arrivals (enqueues) in specified interval: total, by transaction, or by transaction group

Data Type: COUNT (see note on page 121)

Default Title: TRAN ARRIVALS BY TRNCODE

Warning:
```
RM0170W (nn) hh:mm:ss title(parm) = v IN x intrvl
(>thrshld) *****
```

PRCBG - Transactions Processed by Balancing Group (BALG)

Select Code: PRCBG

Parameter: Program (PSB) name

If no parameter is entered, total transactions processed by all BALGs are measured.

Enter program name to measure transactions processed by one BALG.

Enter qualified name to measure transactions processed by a group of BALGs.

Measurement: Transactions processed (dequeues) in specified interval: total, by BALG, or by BALG group

Data Type: COUNT

Default Title: TRANS PROCD BY BALG

Warning:
```
RM0200W (nn) hh:mm:ss title(parm) = v IN x intrvl
(>thrshld) *****
```

PRCCL - Transactions Processed By Class

Select Code: PRCCL

Parameter: Class number

If no parameter is entered, total transactions processed are measured (this includes Fast Path transactions).

Enter a class number to measure transactions processed for one class.

Measurement: Transactions processed (dequeues) in specified interval: total or by class

Data Type: COUNT (see note on page 121)

Default Title: TRANS PROCD BY CLASS

Warning:

```
RM0220W (nn) hh:mm:ss title(parm) = v IN x intrvl
(>thrshld) *****
```

PRCPR - Transactions Processed by Program

Select Code: PRCPR

Parameter: Program (PSB) name

If no parameter is entered, total transactions processed are measured (including Fast Path transactions).

Enter program name to measure transactions processed by one program (including Fast Path transactions).

Enter qualified name to measure transactions processed by a group of programs. If you receive a message that contains the following in the BBI-SS PAS journal:

```
IM4704E ... INVALID PARM - R/C 04
```

the qualified resource could not be found.

Measurement: Transactions processed (dequeues) in specified interval: total, by program, or by program group

Data Type: COUNT (see note on page 121)

Default Title: TRANS PROCD BY PGM

Warning:

```
RM0230W (nn) hh:mm:ss title(parm) = v IN x intrvl
(>thrshld) *****
```

PRCTR - Transactions Processed by Transaction Code

Select Code:	PRCTR
Parameter:	Transaction name If no parameter is entered, total transactions processed are measured (including Fast Path transactions). Enter transaction name to measure the number of times that one transaction was processed. Enter qualified name to measure the number of times a group of transactions was processed. Messages processed in Fast Path regions are not counted (see “PRCBG - Transactions Processed by Balancing Group (BALG)” on page 135). If you receive a message that contains the following in the BBI-SS PAS journal: `IM4704E ... INVALID PARM - R/C 04` a qualified resource could not be found.
Measurement:	Transactions processed (dequeues) in specified interval: total, by transaction, or by transaction group
Data Type:	COUNT (see note on page 121)
Default Title:	TRANS PROCD BY TRANCODE
Warning:	`RM0210W (nn) hh:mm:ss title(parm) = v IN x intrvl (>thrshld) *****`

SCHFL - Scheduling Failures by Type

Select Code:	SCHFL
Parameter:	Failure type If no parameter is entered, total scheduling failures are measured. • Enter INTENT to measure only intent failures • Enter NOOTHER to measure all failures except OTHER • Enter PGM to measure only program conflicts • Enter OTHER to measure only other reasons for failures
Measurement:	Scheduling failures in specified interval: total or by type of failure
Data Type:	COUNT
Default Title:	SCHED FAILURES BY TYPE
Warning:	`RM0240W (nn) hh:mm:ss title(parm) = v IN x intrvl (>thrshld) *****`

Chapter 14. IMS Dependent Region Services

These services measure database and message queue access by programs executing in the dependent regions. For example, to activate the WAIT service to track regions in a long program isolation (PI) wait and issue a warning message if any region is waiting longer than 30 seconds, select the WAIT monitor from the monitor list (SM command) and specify:

```
WMSG ==> MTO
```

```
WVAL ==> 30
```

D2CON - IMS Region Connection to DB2 Subsystem

Select Code: D2CON

Parameter: One- to four-character DB2 subsystem name. If no parameter is entered, all IMS regions connected to a DB2 subsystem are measured.

Measurement: Number of dependent regions connected to DB2

Data Type: STATUS

Default Title: DB2 CONNECTIONS

Warning: `RM0690W (nn) hh:mm:ss title(parm) = v (>thrshld) *****`

D2SON - DB2 Sign On by Subsystem

Select Code: D2SON

Parameter: One- to four-character DB2 subsystem name. If no parameter is entered, all signed-on regions are counted.

Measurement: Number of dependent regions that have issued at least one SQL call and have signed on to DB2

Note: For MPPs or JMPs, the region is counted if at least one transaction has signed on to DB2. This may or may not be the current transaction being processed; for example, an idle region can be counted (see “D2THD - Active IMS Region Threads to DB2 Subsystem” on page 140).

Data Type: STATUS

Default Title: DB2 SIGNON

Warning: `RM1000W DB2 SIGNED-ON REGIONS(TOTAL) = v (>thrshld) *****`

D2THD - Active IMS Region Threads to DB2 Subsystem

Select Code:	D2THD
Parameter:	One- to four-character DB2 subsystem name. If no parameter is entered, all IMS regions that have active threads to a DB2 subsystem are measured.
Measurement:	Number of dependent regions with active threads to DB2
Data Type:	STATUS
Default Title:	DB2 ACTIVE THREADS
Warning:	`RM0700W (nn) hh:mm:ss title(parm) = v (>thrshld) *****`

DBGU - Database Calls per Message Get Unique by Region

Select Code:	DBGU
Parameter:	Region (job) name. (This parameter is required.)
Measurement:	Total DL/I database calls (Fast Path calls not included) performed in specified region since last scheduling divided by number of message get uniques performed (both as measured at time of sample) and compared to warning threshold
Data Type:	Warning Only (threshold must be specified)
Default Title:	DB CALLS/MGU BY RGN
Warning:	`RM0280W (nn) hh:mm:ss title(regname) = v PSB=psbname (>thrshld) *****`

DBTOT - Database Calls per Scheduling by Region

Select Code:	DBTOT
Parameter:	Region (job) name. (This parameter is required.)
Measurement:	Total DL/I database calls (Fast Path calls not included) performed in specified region since last scheduling (measured at time of sample) is compared to the warning threshold
Data Type:	Warning Only (threshold must be specified)
Default Title:	DB CALLS/SCHED BY RGN
Warning:	`RM0270W (nn) hh:mm:ss title(regname) = v PSB=psbname (>thrshld) *****`

DLIDB - DL/I DB Call Time

Select Code:	DLIDB
Parameter:	NOTBMP or region (job) name If no parameter is entered, all regions are monitored. Enter NOTBMP to exclude BMPs or JBPs from consideration by the monitor. Enter a job name to include only that job name for consideration by the monitor.
Measurement:	Number of seconds that the database DL/I call has lasted.
Data Type:	STATUS
Default Title:	DL/I DB CALL TIME
Warning:	`RM1020W (nn) hh:mm:ss title(parm) = value (>thrshld) *****`

Note: You can use the MONDUMPS parameter in BBPARM member IMFBEX00 to specify whether the DLIDB monitor, the DLIDC monitor, or both will take a diagnostic dump when the monitor threshold is met. The diagnostic dumps can be used by IBM or BMC Software support to help diagnose long-runing DLI calls. (See BBPARM member IMFBEX00 for parameter options and recommendations.)

DLIDC - DL/I DC Call Time

Select Code:	DLIDC
Parameter:	GU, GUNOTWFI, NOTWFI, or region (job) name If no parameter is entered, all regions are monitored. Enter GU to monitor only GU calls to the I/O PCB. Enter GUNOTWFI to monitor only GU calls to the I/O PCB for non-WFI and non-PWFI transactions. Enter NOTWFI to monitor all regions except WFI regions. Enter a job name to monitor all transaction DL/I calls for that job name.
Measurement:	Number of seconds that the transaction DL/I call has lasted.
Data Type:	STATUS
Default Title:	DL/I DC CALL TIME
Warning:	`RM1030W (nn) hh:mm:ss title(parm) = value (>thrshld) *****`

Note: You can use the MONDUMPS parameter in BBPARM member IMFBEX00 to specify whether the DLIDB monitor, the DLIDC monitor, or both will take a diagnostic dump when the monitor threshold is met. The diagnostic dumps can be used by IBM or BMC Software support to help diagnose long-runing DLI calls. (See BBPARM member IMFBEX00 for parameter options and recommendations.)

MSGGU - Message Calls per Message Get Unique by Region

Select Code:	MSGGU
Parameter:	Region (job) name. (This parameter is required.)
Measurement:	Total DL/I message queue calls (Fast Path calls not included) performed in specified region since last scheduling divided by number of message get unique performed (both as measured at time of sample) and compared to warning threshold
Data Type:	Warning Only (threshold must be specified)
Default Title:	MSG CALLS/MGU BY RGN
Warning:	`RM0260W (nn) hh:mm:ss title(regname) = v PSB=psbname (>thrshld) *****`

MSGT - Message Calls per Scheduling by Region

Select Code:	MSGT
Parameter:	Region (job) name. (This parameter is required.)
Measurement:	Total DL/I message queue calls (Fast Path calls not included) performed in specified region since last scheduling (as measured at the time of sample) are compared to the warning threshold
Data Type:	Warning Only (threshold must be specified)
Default Title:	MSG CALLS/SCHED BY RGN
Warning:	`RM0250W (nn) hh:mm:ss title(regname) = v PSB=psbname (>thrshld) *****`

WAIT - Region in a Long PI Wait

Select Code: WAIT

Parameter: None

Measurement: A warning is written when any region is detected in a long PI wait (exceeding the threshold value, defined in seconds)

Data Type: Warning Only (threshold must be specified)

Default Title: REGION IN A LONG PI WAIT

Warning: `RM0600W (nn) hh:mm:ss title(TOTAL) = v RGN=regname (>thrshld) *****`

Note: This service can show only one region exception at a time. When the warning condition in one region is cleared, the service continues to check other regions until another exception condition is found and displayed.

Chapter 15. IMS Database Services

These services measure the DL/I buffer pool activity. For example, to activate the DBSTL service to track the number of buffer steal writes in the OSAM database subpool 2 between 10:30 am and 11:40 am, issue a warning message if the count exceeds 100 in a 10 minute interval, and log a plot of the collected data to the BBI-SS PAS Image log at 11:40 when processing of the request stops, select the DBSTL monitor from the monitor list (SM command) and specify:

```
PARM  ==> 2
WVAL  ==> 100
LOG   ==> ATWARN,ATSTOP
START ==> 10:30:00
```

DBIO - Database I/O Count by Subpool

Select Code: DBIO

Parameter: Subpool number

If no parameter is entered, database I/O for the total OSAM pool is measured.

Enter a subpool number to measure database I/O for one subpool.

Measurement: Database I/O in the OSAM pool: total or by subpool

```
= OSAM reads
+ OSAM writes (STEAL)
+ OSAM writes (PURGE)
```

Data Type: COUNT

Default Title: DB I/O COUNT BY SUBPOOL

Warning: `RM0290W (nn) hh:mm:ss title(parm) = v IN x intrvl (>thrshld) *****`

DBHIT - Hit Ratio for OSAM Buffer Pool

Select Code: DBHIT

Parameter: Subpool number

If no parameter is entered, the hit ratio for the entire OSAM buffer pool is measured.

Enter a subpool number to measure hit ratio for one subpool.

Measurement: Hit ratio for the OSAM buffer pool within the interval: total or by subpool

This is a measure of buffer pool performance. A high hit ratio means that a high percentage of reads were satisfied from the buffer pool without having to access external storage.

It is calculated as follows:

```
HIT RATIO = 100 X (A / (A + B))
```

where:

A Number of requests satisfied from the pool

B OSAM reads with I/O

Tuning Tip

Index pools should have a hit ratio of 80.0 or higher. Data pools should be 60.0 or higher.

Data Type: AVERAGE

Default Title: OSAM HIT RATIO BY SP

Warning: `RM0720W (nn) hh:mm:ss title(parm) = v (>thrshld) *****`

DBSTL - Database Buffer Steals by Subpool

Select Code: DBSTL

Parameter: Subpool number

If no parameter is entered, database buffer steal writes for the total OSAM pool are measured.

Enter a subpool number to measure database buffer steal writes for one subpool.

Measurement: Database buffer steal writes in the OSAM pool: total or by subpool

Tuning Tip

Buffer steal writes should be avoided. A high number indicates your buffer pools may be too small.

Data Type: COUNT

Default Title: DB BFR STEALS BY SUBPOOL

Warning: `RM0300W (nn) hh:mm:ss title(parm) = v IN x intrvl (>thrshld) *****`

HPACC - Hiperspace Access by Subpool

Select Code: HPACC

Parameter: Subpool number

If no parameter is entered, all VSAM buffer pools are measured.

Enter a subpool number to measure one subpool.

Measurement: The sum of the successful Hiperspace reads and writes in the interval for the VSAM buffer pool: total or by subpool

Data Type: COUNT

Default Title: H/S ACCESS BY SUBPOOL

Warning: `RM0750W (nn) hh:mm:ss title(parm) = v IN x intrvl (>thrshld) *****`

HPHIT - Hiperspace Hit Ratio by Subpool

Select Code: HPHIT

Parameter: Subpool number

If no parameter is entered, the hit ratio for all VSAM subpools that have associated Hiperspace buffers is calculated.

Enter a subpool number to measure the hit ratio for one subpool.

Measurement: Hiperspace hit ratio within the interval: total or by subpool

The Hiperspace hit ratio indicates how much Hiperspace buffering contributes to the VSAM hit ratio. It is calculated as follows:

```
100 x (A / (B + C))
```

where:

A Number of successful Hiperspace reads for the subpool

B Number of requests satisfied from the VSAM subpool

Note: This number includes the successful Hiperspace reads.

C Number of VSAM reads with I/O

When no parameter is specified, subpools without associated Hiperspace buffers are not included in this calculation. A value of zero is returned if Hiperspace buffering is not supported by IMS.

Data Type: AVERAGE

Default Title: H/S HIT RATIO BY SUBPOOL

Warning: `RM0760W (nn) hh:mm:ss title(parm) = v (>thrshld) *****`

HPSTL - Hiperspace Buffer Steals by Subpool

Select Code:	HPSTL
Parameter:	Subpool number

If no parameter is entered, the number of unsuccessful reads from Hiperspace (buffer steals) for all VSAM subpools is calculated.

Enter a subpool number to measure one subpool.

Measurement: Number of unsuccessful Hiperspace reads: total or by subpool

Each unsuccessful read represents an occasion when data must be obtained from external storage because the Hiperspace pages containing the buffer were stolen. If IMS determines in advance that the desired buffer is not in Hiperspace, it does not even attempt the Hiperspace read. In this case, the number of reads to external storage are greater than the number of unsuccessful Hiperspace reads.

A value of zero is returned if Hiperspace buffering is not supported by IMS.

Tuning Tip

If this number is large, it indicates a shortage of expanded storage that should be investigated.

Data Type:	COUNT
Default Title:	H/S BUFFER STEALS BY SP
Warning:	`RM0810W (nn) hh:mm:ss title(parm) = v IN x intrvl (>thrshld) *****`

SBUSE - Sequential Buffering Storage by Region

Select Code: SBUSE

Parameter: Region name or blank

If no region job or started task name is entered, the sequential buffering storage usage for all regions is calculated.

Enter a region name or qualified name to measure the sequential buffering usage for a region or a group of regions.

Measurement: Kilobytes of virtual storage that IMS currently is using for sequential buffering

Data Type: STATUS

Default Title: SB STORAGE BY REGION

Warning: `RM0730W (nn) hh:mm:ss title(parm) = v (>thrshld) *****`

VDBIO - VSAM Database I/O by Subpool

Select Code: VDBIO

Parameter: Subpool number

If no parameter is entered, database I/O for the total VSAM pool is measured.

Enter a subpool number to measure database I/O for one subpool.

Measurement: Database I/O in the VSAM pool: total or by subpool

```
= reads
+ writes (user-initiated)
+ writes (VSAM-initiated)
```

Data Type: COUNT

Default Title: VSAM DB I/O BY SUBPOOL

Warning: `RM0310W (nn) hh:mm:ss title(parm) = v IN x intrvl (>thrshld) *****`

VDBWR - VSAM Writes by Subpool

Select Code: VDBWR

Parameter: Subpool number

If no parameter is entered, VSAM-initiated writes for the total VSAM pool are measured.

Enter a subpool number to measure VSAM-initiated writes for one subpool.

Measurement: VSAM-initiated writes in the VSAM database pool: total or by subpool

Tuning Tip

VSAM initiated writes should be avoided. A high number indicates your buffer pools may be too small.

Data Type: COUNT

Default Title: VSAM WRITES BY SUBPOOL

Warning: `RM0320W (nn) hh:mm:ss title(parm) = v IN x intrvl (>thrshld) *****`

VHIT - VSAM Hit Ratio by Subpool

Select Code: VHIT

Parameter: Subpool number

If no parameter is entered, the hit ratio for the entire VSAM buffer pool is measured.

Enter a subpool number to measure the hit ratio for one subpool.

Measurement: Hit ratio for the VSAM buffer pool within the interval: total or by subpool

This is a measure of buffer pool performance. A high hit ratio means that a high percentage of reads were satisfied from the buffer pool without having to use external storage.

It is calculated as follows:

```
HIT RATIO = 100 x (A / (A  +  B))
```

where:

A Number of read requests satisfied from the pool

Note: This number includes successful Hiperspace reads.

B Number of reads with I/O

Tuning Tip

Index pools should have a hit ratio of 80.0 or higher. Data pools should be 60.0 or higher.

Data Type: AVERAGE

Default Title: VSAM HIT RATIO BY SP

Warning: `RM0740W (nn) hh:mm:ss title(parm) = v (>thrshld) *****`

Chapter 16. IMS Internals Services

These services measure IMS internal activity for program isolation, pools, logging, latching, and pool allocation. For example, to activate the PSBP service to track the usage of the PSB pool and log a plot of the collected data to the BBI-SS PAS Image log at the end of ten 30-second intervals or every five minutes and issue a warning message if the percentage of the pool usage exceeds 70, select the PSBP monitor from the monitor list (SM command) and specify:

```
INTERVAL ==> 00:00:30
WVAL     ==> 70
LOG      ==> ATPD
```

DBWP - Database Work Area Pool Percentage of Utilization

Select Code:	DBWP
Parameter:	None
Measurement:	DB work area pool percent utilization
	`= bytes allocated * 100 / total pool size in bytes`
Data Type:	PERCENT
Default Title:	DB WORK POOL % UTIL
Warning:	`RM0380W (nn) hh:mm:ss title(TOTAL) = v (>thrshld) *****`

DMBP - DMB Pool Percentage of Utilization

Select Code:	DMBP
Parameter:	None
Measurement:	DMB pool percent utilization
	`= bytes allocated * 100 / total pool size in bytes`
Data Type:	PERCENT
Default Title:	DMB POOL % UTILIZATION
Warning:	`RM0400W (nn) hh:mm:ss title(TOTAL) = v (>thrshld) *****`

DSAP - Dynamic SAP Percentage of Utilization

Select Code:	DSAP
Parameter:	None
Measurement:	Dynamic Save Area pool percent utilization `= number of SAPs assigned * 100 / number of total dynamic SAPs`
Data Type:	PERCENT
Default Title:	DYN SAP % UTILIZATION
Warning:	`RM0330W (nn) hh:mm:ss title(TOTAL) = v (>thrshld) *****`

EPCB - EPCB Pool Percentage of Utilization

Select Code:	EPCB
Parameter:	None
Measurement:	EPCB (Extended PCB) pool percent utilization `= size of EPCB * 100 / total size of EPCB` If Fast Path is not installed in IMS, the percent returned is zero.
Data Type:	PERCENT
Default Title:	EPCB POOL % UTILIZATION
Warning:	`RM0430W (nn) hh:mm:ss title(TOTAL) = v (>thrshld) *****`

LAWT - Average Latch Wait Time

Select Code: LAWT

Parameter: Type, where type can be any of the following latch types for the Common Latch Manager or blank:

ACTL (STATISTICS)
APSB (ALLOC PSB)
BLKM (SMB HASH)
CBTS (CBTS POOL)
CCTL (CCTL for DBCTL latch)
CONV (CONV CHKPT)
DBAU (DBRC AUTH)
DBBP (OSAM BUFP)
DBLK (DEP REGION)
DBLR (BLK LOADER)
DBSL (DB CHKPT)
DCSL (DC CHKPT)
DDRB (DDIR BLK)
DDRP (DDIR POOL)
DISP (DISPATCHER)
DMBB (DMB BLK)
DMBP (DMB POOL)
LOGL (LOG)
LUML (LU 6.2 LUM)
PDRB (PDIR BLK)
PDRP (PDIR POOL)
PSBB (PSB BLK)
PSBP (PSB POOL)
QBSL (QUEUE BUFFER)
QMGR (QUEUE MGR)
SCHD (SCHEDULE)
SMGT (STORAGE MGR)
SUBQ (SCHED SUBQ)
TCTB (TCT BLOCK)
TERM (TERMINAL)
USER (USER)
VLQB (CBTS LQB)
VTCB (CBTS VTCB)
XCNQ (PI ENQ,DEQ)

If no parameter is entered, the average wait time of all latch types is measured.

Measurement: The average WAIT FOR LATCH time for the specified latch type in microseconds within the interval or since the last IMS checkpoint if a checkpoint occurs within the interval.

Data type: AVERAGE

Default Title: AVERAGE LATCH WAIT TIME

Warning:

```
RM0970W (nn) hh:mm:ss title(parm) = v IN x intrvl
(>thrshld) *****
```

LMAWT - Maximum Average Latch Wait Time

Select Code:	LMAWT
Parameter:	None
Measurement:	The latch type with the maximum average wait time is measured If more than one latch type exceeds the threshold, only the one with the largest average wait time is measured.
Data Type:	Warning Only (threshold must be specified) Measurement unit is microseconds.
Default Title:	MAX LATCH WAIT TIME
Warning:	`RM0980W (nn) hh:mm:ss title(parm) = v LATCH=yyyy (>thrshld) *****` where yyyy is the latch type.

MFSP - MFS Pool Percentage of Utilization

Select Code:	MFSP
Parameter:	None

Measurement:	MFS format pool percent utilization
	
	`= bytes allocated * 100`
	`  / dynamic pool size in bytes (available for formats)`

Data Type:	PERCENT
Default Title:	MFS POOL % UTILIZATION
Warning:	`RM0680W (nn) hh:mm:ss title(TOTAL) = v (>thrshld) *****`

OBUFW - OLDS Buffer Waits

Select Code:	OBUFW
Parameter:	None
Measurement:	The number of times within the interval that the IMS logical logger had to wait for a buffer to be written to the OLDS
Data Type:	COUNT
Default Title:	OLDS BUFFER WAITS
Warning:	`RM0910W (nn) hh:mm:ss title(TOTAL) = v IN x intrvl (>thrshld) *****`

OCHKW - OLDS Check Writes

Select Code:	OCHKW
Parameter:	None
Measurement:	The number of Check Write requests to the IMS logical logger within the interval Check Write requests are used mainly for log write ahead for database updates. When IMS encounters a Check Write request, it forces a physical write to the WADS/OLDS if one has not already been done and suspends the requestor.
Data Type:	COUNT
Default Title:	OLDS CHECK WRITES
Warning:	`RM0920W (nn) hh:mm:ss title(TOTAL) = v IN x intrvl (>thrshld) *****`

PIENQ - Program Isolation Enqueues by Region

Select Code:	PIENQ
Parameter:	Region (job) name If no parameter is entered, the total number of PI enqueues held by all dependent regions is measured. Enter region name to measure number of PI enqueues held by one region. Enter qualified name to measure number of PI enqueues held by a group of regions. If you receive a message that contains the following in the BBI-SS PAS journal: `IM4704E ... INVALID PARM - R/C 04` the qualified resource could not be found.
Measurement:	Number of outstanding program isolation enqueues at time of sample: total, by region, or region group
Data Type:	STATUS
Default Title:	PI ENQUEUES BY RGN
Warning:	`RM0630W (nn) hh:mm:ss title(parm) = v (>thrshld) *****`

PIMAX - Maximum Program Isolation Enqueues by Region

Select Code: PIMAX

Parameter: BMP, DBT, IFP, or MPP

If no parameter is entered, all regions are monitored.
Enter:

BMP To monitor only batch message processing (BMP) and Java batch message processing (JBP) regions.

DBT To monitor only DBCTL threads (CICS and ODBA).

IFP To monitor only Fast Path exclusive regions.

MPP To monitor only message processing (MPP) and Java message processing (JMP) regions.

Measurement: The region with the most PI enqueues exceeding the threshold is returned. If more than one region exceeds the threshold, only the one with the most enqueues is returned.

Data Type: Warning Only (threshold must be specified)

Default Title: MAX PI ENQS BY REGION

Warning:

```
RM0710W (nn) hh:mm:ss title(parm) = v xxxxxxxx yyy
(>thrshld) *****
```

where:

xxxxxxxx Is the name of the region that holds the most enqueues

yyy Is the corresponding region number

Note: This service can show only one region exception at a time. When the warning condition clears in a region, the service continues to check other regions until it finds another exception condition.

PIPL - Program Isolation Pool Percentage of Utilization

Select Code: PIPL

Parameter: None

Measurement: Program isolation pool percent utilization, based on its maximum allowable size

```
= number of bytes allocated * 100
/ total number of bytes that can be assigned
```

Data Type: PERCENT

Default Title: PI POOL % UTILIZATION

Warning:

```
RM0640W (nn) hh:mm:ss title(TOTAL) = v (>thrshld) *****
```

POOLA - Pool Allocated Storage

Select Code: POOLA

Parameter: id|CSA|ECSA

where id can be any of the following non-CBT pool IDs:

CIOP
CESS
DBWP
DLDP
DLMP
DPSB
EMHB
EPCB
FPWP
HIOP
LUMC
LUMP
MFBP
PSBW
QBUF
SPAP

CSA Measures the allocated storage in CSA for all pools.

ECSA Measures the allocated storage in ECSA for all pools.

If no parameter is specified, the allocated storage for all pools is measured.

Measurement: The allocated storage in bytes

Data Type: STATUS

Default Title: ALLOCATED POOL STORAGE

Warning: `RM0940W (nn) hh:mm:ss title(parm) = v (>thrshld) *****`

POOLN - Net Expansion Count

Select Code: POOLN

Parameter: id

where id can be any of the following non-CBT fixed pool IDs:

CESS
CIOP
EMHB
FPWP
HIOP
LUMC
LUMP
SPAP

If no parameter is specified, the net expansion for all pools is measured.

Measurement: The difference between total expansions and total compressions of all blocks

Tuning Tip

A high number indicates the primary/secondary allocations are too small. Use the POOLS/DPOOL display services to determine which buffer sizes have the problem.

Data type: STATUS

Default Title: POOL NET EXPANSIONS

Warning: `RM0950W (nn) hh:mm:ss title(parm) = v (>thrshld) *****`

POOLT - Total Expansion/Compression Count

Select Code: POOLT

Parameter: id

where id can be can be any of the following non-CBT fixed pool IDs:

- CESS
- CIOP
- EMHB
- FPWP
- HIOP
- LUMC
- LUMP
- SPAP

If no parameter is specified, the total number of expansions and compressions of blocks for all pools is measured.

Measurement: Total number of expansions and compressions within the interval. The interval needs to be a large value (for example, 01:00:00 for one hour) for the measurement to be meaningful.

Tuning Tip

If net expansion is small as measured by POOLN and total count is high as measured by POOLT, then the blocks are being allocated (expanded) and deleted (compressed) too frequently. It causes extensive IMS overhead, and the primary allocations need to be increased. Use the POOLS/DPOOL display services to determine which buffer sizes have the problem.

Data type: COUNT

Default Title: TOTAL POOL EXPN/COMP

Warning:

```
RM0960W (nn) hh:mm:ss title(parm) = v IN x intrvl
(>thrshld) *****
```

PSBP - PSB Pool Percentage of Utilization

Select Code: PSBP

Parameter: None

Measurement: PSB pool percent utilization

```
= bytes allocated * 100 / total pool size in bytes
```

Note: With the IMS option `LSO=S`, MVIMS monitors DLISAS PSB pool utilization. If the LSO option is not S, total PSB pool utilization is monitored.

Data Type: PERCENT

Default Title: PSB POOL % UTILIZATION

Warning: `RM0390W (nn) hh:mm:ss title(TOTAL) = v (>thrshld) *****`

PSBW - PSB Work Area Pool Percentage of Utilization

Select Code: PSBW

Parameter: None

Measurement: PSB work area pool percent utilization

```
= bytes allocated * 100 / total pool size in bytes
```

Data Type: PERCENT

Default Title: PSB WORK POOL % UTIL

Warning: `RM0370W (nn) hh:mm:ss title(TOTAL) = v (>thrshld) *****`

RECA - RECA Pool Percentage of Utilization

Select Code:	RECA
Parameter:	None
Measurement:	RECA (Receive any) pool percent utilization `# of RECANY buffers in use * 100 / total # of RECANY buffers`
Data Type:	PERCENT
Default Title:	RECA POOL % UTILIZATION
Warning:	`RM0420W (nn) hh:mm:ss title(TOTAL) = v (>thrshld) *****`

WADIO - WADS I/O

Select Code:	WADIO
Parameter:	None
Measurement:	The number of EXCPs to a WADS data set within the interval If dual WADS logging is in effect, the actual number of EXCPs is twice the amount reported.
Data Type:	COUNT
Default Title:	WADS I/O
Warning:	`RM0930W (nn) hh:mm:ss title(TOTAL) = v IN x intrvl (>thrshld) *****`

WKAP - General Work Area Pool Percentage of Utilization

Select Code:	WKAP
Parameter:	None
Measurement:	General work area pool percent utilization `= bytes allocated * 100 / total pool size in bytes`
Data Type:	PERCENT
Default Title:	WKAP % UTILIZATION
Warning:	`RM0360W (nn) hh:mm:ss title(TOTAL) = v (>thrshld) *****`

Chapter 17. OS/390 Services

These services measure the OS/390 resources used by IMS interaction with OS/390. For example, to activate the PAGE service to track the total system paging rates for the IMS control region for 10 sampling intervals and log a plot of the collected data to the BBI-SS PAS Image log when processing of the request completes, select the PAGE monitor from the monitor list (SM command) and specify:

```
PARM ==> IMSCTLRG
LOG  ==> ATSTOP
STOP ==> 10
```

CSAFR - CSA Fragmentation

Select Code:	CSAFR
Parameter:	None
Measurement:	Warning is issued when a storage block of threshold bytes could not be allocated in CSA subpool 241 or unallocated storage at time of sample The threshold should be set equal to the increment specified for PI pool expansion.
Data Type:	Warning Only (threshold must be specified)
Default Title:	CSA FRAGMENTATION
Warning:	`RM0480W (nn) hh:mm:ss title(TOTAL) = v (<thrshld) *****`

CSAUT - CSA Percentage of Utilization

Select Code:	CSAUT
Parameter:	None
Measurement:	CSA percent utilization at time of sample (virtual) `= bytes of CSA allocated * 100 / total size of CSA in bytes`
Data Type:	PERCENT
Default Title:	CSA % UTILIZATION
Warning:	`RM0460W (nn) hh:mm:ss title(TOTAL) = v (>thrshld) *****`

DLIO - DL/I EXCP Count by ddname

Select Code:	DLIO
Parameter:	ddname
	If no parameter is specified, total EXCPs for all data sets allocated to the IMS DLISAS region are measured.
	Enter a ddname to measure EXCPs for one data set allocated to IMS.
	Enter qualified name to measure EXCPs for a group of data sets.
Measurement:	Number of successful EXCPs for DL/I data sets allocated to the DLISAS region (`LSO=S`): total or by ddname
	If any other LSO option is in effect, this service measures data sets in the control region, the same as SYSIO.
	If SMF is inactive, the count will be zero for non-VSAM and non-OSAM data sets.
Data Type:	COUNT
Default Title:	DL/I EXCP COUNT BY DDNAME
Warning:	`RM0510W (nn) hh:mm:ss title(parm) = v IN x intrvl (>thrshld) *****`

DPAGE - Demand Paging by Region

Select Code:	DPAGE
Parameter:	Region (job) name
	If no parameter is entered, total demand paging for the system is measured.
	Enter region name (IMS regions only) to measure demand paging for one region.
Measurement:	Total page-ins in the specified interval: total or by region (includes reclaims)
Data Type:	COUNT
Default Title:	DEMAND PAGING BY REGION
Warning:	`RM0450W (nn) hh:mm:ss title(parm) = v IN x intrvl (>thrshld) *****`

ECSAU - Extended CSA Percent Utilization

Select Code: ECSAU

Parameter: None

Measurement: Extended CSA percent utilization at the time of sample

This is computed as:

```
100 x (A - B) / A
```

where:

`A` Total ECSA

`B` Space described by FBQES in ECSA

Note: Using this formula, a 4K block of ECSA is recognized as 100% utilized, even if only a portion is actually allocated.

Data Type: PERCENT

Default Title: ECSA % UTILIZATION

Warning: `RM0800W (nn) hh:mm:ss title(TOTAL) = v (>thrshld) *****`

PAGE - Paging (Region)

Select Code: PAGE

Parameter: Region (job) name

If no parameter is entered, total paging for the system is measured.

Enter region name (IMS regions only) to measure paging for one region.

Measurement Total private area, VIO, and swap page-ins and page-outs in the specified interval: total or by region (does not include reclaims)

Data Type: COUNT

Default Title: PAGING (REGION)

Warning: `RM0440W (nn) hh:mm:ss title(parm) = v IN x intrvl (>thrshld) *****`

SYSIO - EXCP Count by ddname

Select Code:	SYSIO
Parameter:	ddname
	If no parameter is specified, total EXCPs for all data sets allocated to the IMS control region are measured.
	Enter a ddname to measure EXCPs for one data set allocated to IMS.
	Enter a qualified data set name to measure EXCPs for a group of data sets.
Measurement:	Number of successful EXCPs: total or by ddname
	If SMF is inactive, the count will be zero for non-VSAM and non-OSAM data sets.
Data Type:	COUNT
Default Title:	EXCP COUNT BY DDNAME
Warning:	`RM0530W (nn) hh:mm:ss title(parm) = v IN x intrvl (>thrshld) *****`

Chapter 18. IRLM Services

These services measure critical IRLM activity when IMS IRLM is used. For example, to activate the LHELD service every minute to monitor for the number of locks held by IRLM for a target IMS and issue a warning message when it holds one or more, select the LHELD monitor from the monitor list (SM command) and specify:

```
INTERVAL ==> 00:01:00
WVAL     ==> NZ
```

LDLCK - Number of Deadlocks

Select Code:	LDLCK
Parameter:	None
Measurement:	Number of deadlocks that occurred within each user-specified interval.
Data Type:	COUNT
Default Title:	IRLM DEADLOCKS
Warning:	`RM0820W (nn) hh:mm:ss title(parm) = v IN x intrvl (>thrshld) *****`

LHELD - Number of Locks Held

Select Code:	LHELD
Parameter:	IMS online region number If no parameter is entered, all the locks being held by the target IMS are counted. Enter IMS online region number to measure the locks being held by that region.
Measurement:	Number of locks currently being held by an IMS online dependent region or by the target IMS as a whole
Data Type:	STATUS
Default Title:	IRLM LOCKS HELD
Warning:	`RM0890W (nn) hh:mm:ss title(parm) = v (>thrshld) *****`

LKMAX - Maximum Locks Held by Region

Select Code: LKMAX

Parameter: BMP, DBT IFP, or MPP

If no parameter is entered, all regions are monitored..
Enter:

BMP	To monitor only batch message processing (BMP) and Java batch message processing (JBP) regions.
DBT	To monitor only DBCTL threads (CICS and ODBA).
IFP	To monitor only Fast Path exclusive regions.
MPP	To monitor only message processing (MPP) and Java message processing (JMP) regions.

Measurement: Region with the most outstanding IRLM locks over the threshold is returned

If more than one region is over the threshold, the one with the most locks is returned.

Data Type: Warning Only (threshold must be specified)

Default Title: MAX LOCKS HELD BY REGION

Warning:

```
RM0870W (nn) hh:mm:ss title(parm) = v xxxxxxxx yyy
(>thrshld) *****
```

where:

`xxxxxxxx`	Is the region name
`yyy`	Is the region number

Note: This service shows only one online dependent region exception at a time. When the warning condition in the region clears, it continues to check other online dependent regions until another exception is found.

LKREQ - Number of Lock Requests

Select Code:	LKREQ
Parameter:	GLOBAL If no parameter is entered, the subset related to the target IMS is measured. Enter GLOBAL to measure the entire IRLM system.
Measurement:	Number of lock requests
Data Type:	COUNT
Default Title:	IRLM LOCK REQUESTS
Warning:	`RM0840W (nn) hh:mm:ss title(parm) = v IN x intrvl (>thrshld) *****`

LSUSP - Number of Suspensions

Select Code:	LSUSP
Parameter:	None
Measurement:	Number of IRLM suspensions that occurred within each user-specified interval
Data Type:	COUNT
Default Title:	IRLM SUSPENSIONS
Warning:	`RM0900W (nn) hh:mm:ss title(parm) = v IN x intrvl (>thrshld) *****`

LWAIT - Region in IRLM Suspend

Select Code: LWAIT

Parameter: BMP, DBT, IFP, or MPP

If no parameter is entered, all regions are monitored.
Enter:

BMP To monitor only batch message processing (BMP) and Java batch message processing (JBP) regions.

DBT To monitor only DBCTL threads (CICS and ODBA), valid only for IRLM 2.1 and later.

IFP To monitor only Fast Path exclusive regions.

MPP To monitor only message processing (MPP) and Java message processing (JMP) regions.

Measurement: A warning is written when a monitored region is in IRLM suspend state longer than the threshold value specified in seconds

Data Type: Warning Only (threshold must be specified)

Default Title: REGION IN IRLM SUSPEND

Warning:

```
RM0880W (nn) hh:mm:ss title(parm) = v xxxxxxxx yyy
(>thrshld) *****
```

where:

`xxxxxxxx` Is the region name
`yyy` Is the region number

Note: This service shows only one online dependent region exception at a time. When the warning condition in the region clears, it continues to check other online dependent regions until another exception is found.

LWNUM - Number of Suspended IRLM Requests

Select Code: LWNUM (IRLM 2.1 and later)

Parameter: GLOBAL, GBL, GBL`ssss`, or `ssss`

If no parameter is entered, all suspended requests from the target IMS are measured. Enter:

GLOBAL or GBL
: All suspended requests are counted (not limited to the target IMS).

GBL`ssss`
: Where `ssss` is number of seconds. All requests suspended longer than `ssss` seconds are counted.

`ssss`
: Where `ssss` is number of seconds. All requests from the target IMS suspended longer than `ssss` seconds are counted.

Measurement: Number of IRLM requests suspended at the end of each interval that have lasted longer than a user-specified number of seconds

Data Type: STATUS

Default Title: # REGIONS IN IRLM SUSPND

Warning: `RM1010W (nn) hh:mm:ss title(parm) = v (>thrshld) *****`

PTBLK - Number of PTB Locks

Note: This service applies to IRLM 1.5.

Select Code: PTBLK

Parameter: GLOBAL

If no parameter is entered, the subset related to the target IMS is measured.

Enter GLOBAL to measure the entire IRLM system.

Measurement: Number of lock requests that resulted in PTB process

The monitor is not valid if the count is not meaningful. This count is only meaningful for an IRLM started with parameter `SCOPE=GLOBAL`.

Data Type: COUNT

Default Title: IRLM PTB LOCKS

Warning: `RM0850W (nn) hh:mm:ss title(parm) = v IN x intrvl (>thrshld) *****`

VSEND - Number of VTAM Sends

Note: This service applies to IRLM 1.5.

Select Code: VSEND

Parameter: None

Measurement: Number of VTAM sends issued by the IRLM system

Use this monitor only if you specify SCOPE=GLOBAL in your IRLM procedure. If SCOPE=LOCAL is specified, VTAM is not used by IRLM. VSEND is only for monitoring VTAM sends.

Data Type: COUNT

Default Title: IRLM VTAM SENDS

Warning: `RM0860W (nn) hh:mm:ss title(parm) = v IN x intrvl (>thrshld) *****`

Chapter 19. Workload Monitor Services (Quick Reference)

This section provides an alphabetical list of all the Workload Monitor services and their parameters with page references to a more detailed description about their use. See Chapter 8, “Displaying a List of Data Collection Monitors (SM Command)” on page 87 for a complete description of the Data Collection Monitors list application.

Table 7. Workload Monitor Service Select Codes

Service Select Code	**Parameter**	**See**
$CBMP	[id],[keywords]	“$CBMP - BMP and JBP Region Calls” on page 187
$CMPP	[id],[keywords]	“$CMPP - MPP and JMP Region Calls” on page 189
$CTOT	[id],[keywords]	“$CTOT - All Region Calls” on page 190
$CDBT	[id],[keywords]	“$CDBT - DBCTL Region DLI Calls” on page 188
#CDB2	[id],[keywords]	“#CDB2 - DB2 Data Access Calls” on page 179
#CIC	[id],[keywords]	“#CIC - Control Interval Contentions” on page 183
#OBAW	[id],[keywords]	“#OBAW - OBA Latch Waits” on page 184
#PROC	[id],[keywords]	“#PROC - Number of Transactions Processed” on page 197
#SDB2	[id],[keywords]	“#SDB2 - DB2 Nondata Access Calls” on page 180
@ELAP	[id],[keywords]	“@ELAP - Average Elapsed Time” on page 191
@INPQ	[id],[keywords]	“@INPQ - Average Input Queue Time” on page 193
@OBA	[id],[keywords]	“@OBA - Overflow Buffer Usage” on page 185
@PDB2	[id],[keywords]	“@PDB2 - DB2 CPU Time” on page 181
@RESP	[id],[keywords]	“@RESP - Average Response Time” on page 194

Chapter 20. Workload Monitor Support for DBCTL and ODBA Threads

This chapter summarizes service support for DBCTL threads.

Collecting DBCTL Thread Data

To collect DBCTL thread data, CICS=YES or CICS=ONLINE must be specified in BBPARM member IMFECP00. Then, when you request a monitor service (SM primary command), you can include or exclude the DBCTL thread data (CICS and ODBA) with the PGMTYPE workload selection keyword as follows:

- PGMTYPE=DBT collects just DBCTL thread data.
- PGMTYPE=NOTDBT excludes DBCTL thread data collection from a Workload Monitor request.

DBCTL thread data support is summarized below.

Service Support

The following Workload Monitor services provide DBCTL thread data:

#PROC
@ELAP
@OBA
#OBAW
#CIC
$CTOT
$CDBT

If you use @ELAP or $CDBT, you should consider the following:

@ELAP You may want to exclude DBCTL thread transactions from @ELAP monitoring by specifying PGMTYPE=NOTDBT.

The elapsed time for conversational DBCTL thread transactions may be quite high because it includes user think time at the terminal. The Event Collector measures elapsed time from transaction start to the next SYNC point. With a conversational DBCTL thread transaction, SYNC point may not occur until transaction termination.

If a transaction issues more than one commit or allocates more than one PSB per execution, the elapsed time reflects only the time between SYNC points, not the elapsed time of the DBCTL thread transaction.

If you want to monitor conversational and nonconversational DBCTL thread transactions separately, you must use the PSB and/or TRAN workload selection criteria keyword when you request a monitor service.

$CDBT The $CDBT monitor shows the number of DLI calls by DBCTL threads and complements the $CTOT, $CBMP, and $CMPP monitors. You can specify a selection criteria of CALLTYPE=ALL or DLI with the monitor request. Specifying CALLTYPE=MSG or DB2 results in zero counts.

The following Workload Monitor services do not provide DBCTL data because DB2 is not accessible through the DBCTL interface.

@PDB2
#CDB2
#SDB2

The following Workload Monitor services do not provide DBCTL data because DBCTL transactions do not go through the IMS message queue.

@INPQ
@RESP

Chapter 21. DB2 Activity Monitors

These services measure the amount of dependent region CPU time used by a transaction to make DB2 requests and the number of DB2 calls made through the IMS Attach facility.

#CDB2 - DB2 Data Access Calls

Select Code: #CDB2

Parameter: User-defined identifier for the request or blank.

Measurement: A count of the number of SQL calls to DB2 for data access to tables for the selected workload and sampled interval. If workload selection is not requested, all calls to DB2 are counted.

Note: DBCTL data are not collected by this service because DB2 is not accessible through the DBCTL interface.

Data Type: COUNT

Selections: Valid workload selection keyword parameters:

```
TRANTYPE    ==> DB2
TRAN        ==> trncode
CLASS       ==> nn
PROG        ==> name
PSB         ==> name
REGION      ==> jobname
RGNID       ==> nnn
TERM        ==> name|SYNCLOCK|MSCCLOCK
USERID      ==> userid
PGMTYPE     ==> MPP|MDP|IFP|FPU|TPI|BMP|NOTDBT
```

The MPP program type includes the JMP program type, the BMP type includes the JBP type, and NOTDBT includes ODBA threads.

All keyword operands are ORed together. Multiple workload selection keyword parameters are ANDed. The keyword parameters and operands are described in Table 4 on page 101.

Note: TRANTYPE is set internally to DB2 to ensure correct averaging by excluding non-DB2 transactions.

Default Title: DB2 DATABASE CALLS

Warning: `WM0090W (nn) hh:mm:ss title(parm) = v IN x intrvl (>thrshld) *****`

Example

```
PARM  ==> PROFILE1
START ==> 00:11:00
STOP  ==> 00:12:00
LOG   ==> ATPD,ATSTOP
```

#SDB2 - DB2 Nondata Access Calls

Select Code: #SDB2

Parameter: User-defined identifier for the request or blank.

Measurement: A count of the number of nondata-access SQL calls (DDL, dynamic, and control-type calls) to DB2 for the selected workload and sampled interval. If workload selection is not requested, all nondata access calls to DB2 are counted.

Note: DBCTL data is not collected by this service because DB2 is not accessible through the DBCTL interface.

Data Type: COUNT

Selections: Valid workload selection keyword parameters:

```
TRANTYPE    ==> DB2
TRAN        ==> trncode
CLASS       ==> nn
PROG        ==> name
PSB         ==> name
REGION      ==> jobname
RGNID       ==> nnn
TERM        ==> name|SYNCLOCK|MSCCLOCK
USERID      ==> userid
PGMTYPE     ==> MPP|MDP|IFP|FPU|TPI|BMP|NOTDBT
```

The MPP program type includes the JMP program type, the BMP type includes the JBP type, and NOTDBT includes ODBA threads.

All keyword operands are ORed together. Multiple workload selection keyword parameters are ANDed. The keyword parameters and operands are described in Table 4 on page 101.

Note: TRANTYPE is set internally to DB2 to ensure correct averaging by excluding non-DB2 transactions.

Default Title: DB2 SPECIAL CALLS

Warning: `WM0100W (nn) hh:mm:ss title(parm) = v IN x intrvl (>thrshld) *****`

Example

```
PARM  ==> NONO
WVAL  ==> 1
TITLE ==> DISALLOWED DB2 CALLS
```

@PDB2 - DB2 CPU Time

Select Code: @PDB2

Parameter: User-defined identifier for the request or blank.

Measurement: The average CPU time used per transaction to process DB2 requests for the selected workload and sampled interval. If workload selection is not requested, the DB2 CPU time for the total workload is measured.

Note: DBCTL data is not collected by this service because DB2 is not accessible through the DBCTL interface.

Data Type: AVERAGE

Selections: Valid workload selection keyword parameters:

```
TRANTYPE    ==> DB2
TRAN        ==> trncode
CLASS       ==> nn
PROG        ==> name
PSB         ==> name
REGION      ==> jobname
RGNID       ==> nnn
TERM        ==> name|SYNCLOCK|MSCCLOCK
USERID      ==> userid
PGMTYPE     ==> MPP|MDP|IFP|FPU|TPI|BMP|NOTDBT
```

The MPP program type includes the JMP program type, the BMP type includes the JBP type, and NOTDBT includes ODBA threads.

All keyword operands are ORed together. Multiple workload selection keyword parameters are ANDed. The keyword parameters and operands are described in Table 4 on page 101.

Note: TRANTYPE is set internally to DB2 to ensure correct averaging by excluding non-DB2 transactions.

Default Title: AVG DB2 CPU TIME

Warning: `WM0080W (nn) hh:mm:ss title(parm) = v (>thrshld) *****`

Example

```
PARM   ==> ACCTG
USERID ==> ACC+
```

Chapter 22. Fast Path Activity Monitors

These services measure the use of IMS Fast Path resources, which include CI contention, waits for OBA, and overflow buffer usage. They can also be used for DBCTL threads.

#CIC - Control Interval Contentions

Select Code: #CIC

Parameter: User-defined identifier for the request or blank.

Measurement: The number of Fast Path database control interval contentions for the selected workload and sampled interval. If workload selection is not requested, the control interval contentions for the total workload measured.

Data Type: COUNT

Selections: Valid workload selection keyword parameters:

```
TRANTYPE    ==> FP
TRAN        ==> trncode
CLASS       ==> nn
PROG        ==> name
PSB         ==> name
REGION      ==> jobname
RGNID       ==> nnn
TERM        ==> name|SYNCLOCK|MSCCLOCK
USERID      ==> userid
PGMTYPE     ==> MPP|MDP|IFP|FPU|TPI|BMP|DBT|NOTDBT
```

The MPP program type includes the JMP program type, the BMP type includes the JBP type, and NOTDBT includes ODBA threads.

All keyword operands are ORed together. Multiple workload selection keyword parameters are ANDed. The keyword parameters and operands are described in Table 4 on page 101.

Note: TRANTYPE is set internally to FP to ensure correct averaging by excluding transactions that do not access Fast Path databases.

Default Title: CI CONTENTIONS

Warning: `WM0070W (nn) hh:mm:ss title(parm) = v IN x intrvl (>thrshld) *****`

Example

```
WVAL ==> 3
WMSG ==> MTO
WLIM ==> 10
LOG  ==> ATWARN
```

#OBAW - OBA Latch Waits

Select Code: #OBAW

Parameter: User-defined identifier for the request or blank.

Measurement: The number of transaction OBA latch waits for the selected workload and sampled interval. If workload selection is not requested, the total OBA latch waits by Fast Path transactions is measured.

Data Type: COUNT

Selections: Valid workload selection keyword parameters:

```
TRANTYPE    ==> FP
TRAN        ==> trncode
CLASS       ==> nn
PROG        ==> name
PSB         ==> name
REGION      ==> jobname
RGNID       ==> nnn
TERM        ==> name|SYNCLOCK|MSCCLOCK
USERID      ==> userid
PGMTYPE     ==> MPP|MDP|IFP|FPU|TPI|BMP|DBT|NOTDBT
```

The MPP program type includes the JMP program type, the BMP type includes the JBP type, and NOTDBT includes ODBA threads.

All keyword operands are ORed together. Multiple workload selection keyword parameters are ANDed. The keyword parameters and operands are described in Table 4 on page 101.

Note: TRANTYPE is set internally to FP to ensure correct averaging by excluding transactions that do not access Fast Path databases.

Default Title: OVERFLOW BUFFER WAITS

Warning: `WM0060W(nn) hh:mm:ss title(parm) = v IN x intrvl (>thrshld) *****`

Example

```
REGION ==> PROD+
```

@OBA - Overflow Buffer Usage

Select Code: @OBA

Parameter: User-defined identifier for the request or blank.

Measurement: The average overflow buffer usage by Fast Path transactions for the selected workload and sampled interval. If workload selection is not requested, the overflow buffer usage by Fast Path transactions for the total workload is measured.

Data Type: AVERAGE

Selections: Valid workload selection keyword parameters:

```
TRANTYPE    ==> FP
TRAN        ==> trncode
CLASS       ==> nn
PROG        ==> name
PSB         ==> name
REGION      ==> jobname
RGNID       ==> nnn
TERM        ==> name|SYNCLOCK|MSCCLOCK
USERID      ==> userid
PGMTYPE     ==> MPP|MDP|IFP|FPU|TPI|BMP|DBT|NOTDBT
```

The MPP program type includes the JMP program type, the BMP type includes the JBP type, and NOTDBT includes ODBA threads.

All keyword operands are ORed together. Multiple workload selection keyword parameters are ANDed. The keyword parameters and operands are described in Table 4 on page 101.

Note: TRANTYPE is set internally to FP to ensure correct averaging by excluding transactions that do not access Fast Path databases.

Default Title: AVG OVFLW ALLOCATION

Warning: `WM0050W (nn) hh:mm:ss title(parm) = v (>thrshld) *****`

Example

```
PARM ==> WHSE
PSB  ==> WH+
```

Chapter 23. Global Region Call Monitors

The global region call monitors measure all region calls (MPP, JMP, IFP, BMP, and JBP) and DLI calls issued by DBCTL threads ($CTOT and $CDBT services) as they occur during IMS application program processing. Call types by region can be selected with the use of the CALLTYPE keyword specified with the service request.

$CBMP - BMP and JBP Region Calls

Select Code:	$CBMP
Parameter:	User-defined identifier for the request or blank.
Measurement:	A count of the calls issued by BMP and JBP regions in the sampled interval.
Data Type:	COUNT
Selections:	Valid call type selection keyword parameters: `CALLTYPE ==> ALL|MSG|DB2|DLI` where ALL is for all database and message queue calls; MSG is only message queue calls; DB2 is only DB2 calls; and DLI is only DL/I calls. If CALLTYPE is not specified, the default is ALL. Multiple operands are ORed together.
Default Title:	TOTAL BMP CALLS
Warning:	`WM0120W (nn) hh:mm:ss title(parm) = v IN x intrvl (>thrshld) *****`

Example

```
PARM     ==> DLICALLS
TITLE    ==> BMP DL/I CALLS
CALLTYPE ==> DLI
```

$CDBT - DBCTL Region DLI Calls

Select Code: $CDBT

Parameter: User-defined identifier for the request or blank.

Measurement: A count of the DL/I calls issued by DBCTL threads in the sampled interval. The count is a total of all calls or by call type as specified with the request.

Data Type: COUNT

Selections: Valid call type selection keyword parameters:

```
CALLTYPE  ==> ALL|DLI
```

Note: IMS processes only DLI calls for this region type; therefore, specifying `CALLTYPE ==> MSG` or `CALLTYPE ==> DB2` results in zero counts.

Default Title: TOTAL DBT (CICS) CALLS

Warning:

```
WM0140W (nn) hh:mm:ss title(parm) = v IN x intrvl
(>thrshld) *****
```

Example

```
PARM     ==> DLIDBT1
WVAL     ==> 100
CALLTYPE ==> DLI
```

$CMPP - MPP and JMP Region Calls

Select Code: $CMPP

Parameter: User-defined identifier for the request or blank.

Measurement: A count of the calls issued by MPP, JMP, and IFP regions in the sampled interval. The count is a total of all calls or by call type as specified with the request.

Data Type: COUNT

Selections: Valid call type selection keyword parameters:

```
CALLTYPE ==> ALL|MSG|DB2|DLI
```

where ALL is for all database and message queue calls; MSG is only message queue calls; DB2 is only DB2 calls; and DLI is only DL/I calls. The default is ALL. Multiple operands are ORed together.

Default Title: TOTAL MPP CALLS

Warning:
```
 WM0130W (nn) hh:mm:ss title(parm) = v IN x intrvl
(>thrshld) *****
```

Example

```
PARM     ==> DB2CALLS
TITLE    ==> MPP IFP DB2 CALLS
CALLTYPE ==> DB2
```

$CTOT - All Region Calls

Select Code: $CTOT

Parameter: User-defined identifier for the request or blank.

Measurement: A count of the calls issued by all region types in the sampled interval. The count is a total of all calls or by call type as specified with the request.

Data Type: COUNT

Selections: Valid call type selection keyword parameters:

```
CALLTYPE ==> ALL|MSG|DB2|DLI
```

where ALL is for all database and message queue calls; MSG is only message queue calls; DB2 is only DB2 calls; and DLI is only DL/I calls. The default is ALL. Multiple operands are ORed together.

Default Title: TOTAL PROGRAM CALLS

Warning:

```
WM0110W (nn) hh:mm:ss title(parm) = v IN x intrvl
(>thrshld) *****
```

Example

```
PARM     ==> DB2CALLS
CALLTYPE ==> DB2
```

Chapter 24. Elapsed Timing Monitors

These services measure the transaction input queue time, elapsed transaction execution time (including elapsed time for DBCTL threads), transaction response time in seconds for MPP or IFP transactions, and transaction response time to the input terminal.

@ELAP - Average Elapsed Time

Select Code: @ELAP

Parameter: User-defined identifier for the request or blank.

Measurement: Average MPP and JMP transaction, IFP transaction, or DBCTL thread elapsed execution time for the selected workload and sampled interval. If workload selection is not requested, the average elapsed execution time for the total workload is measured.

Data Type: AVERAGE

Selections: Valid workload selection keyword parameters:

```
TRANTYPE    ==> ALL|DLI|DB2|FP
TRAN        ==> trncode
CLASS       ==> nn
PROG        ==> name
PSB         ==> name
REGION      ==> jobname
RGNID       ==> nnn
TERM        ==> name|SYNCLOCK|MSCCLOCK
USERID      ==> userid
PGMTYPE     ==> MPP|MDP|IFP|FPU|TPI|BMP|DBT|NOTDBT
```

The MPP program type includes the JMP program type, the BMP type includes the JBP type, and NOTDBT includes ODBA threads.

All keyword operands are ORed together. Multiple workload selection keyword parameters are ANDed. The keyword parameters and operands are described in Table 4 on page 101.

Note: You may want to exclude DBCTL thread transactions by specifying `PGMTYPE=NOTDBT`. The elapsed time for conversational monitors may be quite high because it includes user think time at the terminal. Elapsed time is measured from transaction start to SYNC point. SYNC point for a conversational DBCTL thread transaction may not occur until the transaction terminates. To monitor conversational and nonconversational DBCTL thread transactions separately, use the PSB and TRAN workload selection keyword parameters.

If a DBCTL thread transaction issues more than one commit or allocates more than one PSB per execution, the elapsed time reflects only the time between SYNC points, not the elapsed time of the transaction.

Default Title: AVG ELAPSED TIME

Warning: `WM0020W (nn) hh:mm:ss title(parm) = v (>thrshld) *****`

Example

```
PARM  ==> CLS01
WVAL  ==> 2.50
TITLE ==> CLASS 1 ELAPSED TIME
CLASS ==> 1
```

@INPQ - Average Input Queue Time

Select Code: @INPQ

Parameter: User-defined identifier for the request or blank.

Measurement: Average MPP and JMP or IFP transaction input queue time for the selected workload and sampled interval. If workload selection is not requested, the average input queue time for the total workload is measured.

Note: DBCTL thread transaction data is not collected, because these transactions do not go through the IMS message queue.

Data Type: AVERAGE

Selections: Valid workload selection keyword parameters:

```
TRANTYPE    ==> ALL|DLI|DB2|FP
TRAN        ==> trncode
CLASS       ==> nn
PROG        ==> name
PSB         ==> name
REGION      ==> jobname
RGNID       ==> nnn
TERM        ==> name|SYNCLOCK|MSCCLOCK
USERID      ==> userid
PGMTYPE     ==> MPP|MDP|IFP|FPU|TPI|BMP|DBT|NOTDBT
```

The MPP program type includes the JMP program type, the BMP type includes the JBP type, and NOTDBT includes ODBA threads.

All keyword operands are ORed together. Multiple workload selection keyword parameters are ANDed. The keyword parameters and operands are described in Table 4 on page 101.

Default Title: AVG INPUT Q TIME

Warning: `WM0010W (nn) hh:mm:ss title(parm) = v (>thrshld) *****`

Example

```
PARM ==> R1
LOG  ==> ATPD
TRAN ==> A++21+
```

@RESP - Average Response Time

Select Code: @RESP

Parameter: User-defined identifier for the request or blank.

Measurement: Average MPP and JMP or IFP transaction response time (input queue time plus elapsed time) for the selected workload and sampled interval.
If workload selection is not requested, the average response time for the total workload is measured.

Note: DBCTL thread transaction data is not collected, because these transactions do not go through the IMS message queue.

Data Type: AVERAGE

Selections: Valid workload selection keyword parameters:

```
TRANTYPE    ==> ALL|DLI|DB2|FP
TRAN        ==> trncode
CLASS       ==> nn
PROG        ==> name
PSB         ==> name
REGION      ==> jobname
RGNID       ==> nnn
TERM        ==> name|SYNCLOCK|MSCCLOCK
USERID      ==> userid
PGMTYPE     ==> MPP|MDP|IFP|FPU|TPI|BMP|DBT|NOTDBT
```

The MPP program type includes the JMP program type, the BMP type includes the JBP type, and NOTDBT includes ODBA threads.

All keyword operands are ORed together. Multiple workload selection keyword parameters are ANDed. The keyword parameters and operands are described in Table 4 on page 101.

Default Title: AVG RESPONSE TIME

Warning: `WM0030W (nn) hh:mm:ss title(parm) = v (>thrshld) *****`

Example

```
PARM     ==> MDCLS
INTERVAL ==> 00:03:00
WVAL     ==> 6
RANGES   ==> 3,6,9,12
LOG      ==> ATPD,ATSTOP
CLASS    ==> 4,5
```

@TRSP - Average Terminal Response Time

Select Code: @TRSP

Parameter: User-defined identifier for the request or blank.

Measurement: Average time between message arrival to IMS and the enqueue of the response to the input terminal. Response time is calculated as the difference between the time the original input message was received by IMS and the commit time of the transaction that inserted the response.

Fast Path transactions are not included, and the time between commit and reception of the message at the terminal is not included.

The reported average response time may be inaccurate if

- the response was sent prior to commit time, by way of an express mode PCB for example
- PGM-to-PGM switches occurred and more than one application responded to the input terminal

Note: DBCTL thread transaction data is not collected because these transactions do not go through the IMS message queue.

Data Type: AVERAGE

Selections: Valid workload selection keyword parameters:

```
TRANTYPE    ==> ALL|DLI|DB2|FP
TRAN        ==> trncode
CLASS       ==> nn
PROG        ==> name
PSB         ==> name
REGION      ==> jobname
RGNID       ==> nnn
TERM        ==> name|SYNCLOCK|MSCCLOCK
USERID      ==> userid
PGMTYPE     ==> MPP|MDP|IFP|FPU|TPI|BMP|DBT|NOTDBT
```

The MPP program type includes the JMP program type, the BMP type includes the JBP type, and NOTDBT includes ODBA threads.

All keyword operands are ORed together. Multiple workload selection keyword parameters are ANDed. The keyword parameters and operands are described in Table 4 on page 101.

Default Title: AVG TERMINAL RESPONSE

Warning: `WM0030W (nn) hh:mm:ss title(parm) = v (>thrshld) *****`

Example

```
PARM     ==> MDCLS
INTERVAL ==> 00:03:00
```

```
WVAL      ==> 6
RANGES    ==> 3,6,9,12
LOG       ==> ATPD,ATSTOP
CLASS     ==> 4,5
```

Chapter 25. Transaction Monitor

This monitor measures the number of transactions processed by IMS.

#PROC - Number of Transactions Processed

Select Code: #PROC

Parameter: User-defined identifier for the request or blank.

Measurement: A count of the MPP and JMP transactions or IFP transactions processed for the selected workload and sampled interval. If workload selection is not requested, all the transactions are counted.

Data Type: COUNT

Selections: Valid workload selection keyword parameters:

```
TRANTYPE    ==> ALL|DLI|DB2|FP
TRAN        ==> trncode
CLASS       ==> nn
PROG        ==> name
PSB         ==> name
REGION      ==> jobname
RGNID       ==> nnn
TERM        ==> name|SYNCLOCK|MSCCLOCK
USERID      ==> userid
PGMTYPE     ==> MPP|MDP|IFP|FPU|TPI|BMP|DBT|NOTDBT
```

The MPP program type includes the JMP program type, the BMP type includes the JBP type, and NOTDBT includes ODBA threads.

All keyword operands are ORed together. Multiple workload selection keyword parameters are ANDed. The keyword parameters and operands are described in Table 4 on page 101.

Default Title: TRANS PROCESSED

Warning: `WM0040W (nn) hh:mm:ss title(parm) = v IN x intrvl (>thrshld) *****`

Example 1

```
PARM ==> ALL
```

Example 2

```
PARM     ==> ALLDB2
TRANTYPE ==> DB2
```

Part 4. Traces

The trace services are used to activate and display a summary workload trace or detail transaction processing event trace.

Chapter 26. Tracing Transactions

Workload Analyzer provides unique application tracing that allows you to selectively trace a transaction and view how well your application is performing with integrated trace displays. Application tracing consists of data collection and display services that provide:

Summary trace
: Response time of the transaction, the application CPU time, the IMS CPU time, the number of database calls (DL/I and DB2), and the number of I/Os.

Detail trace
: Each type of database call, its return code and elapsed time (in the order processed), the name of the database accessed, and associated database I/O data and segment search argument, key feedback, and I/O area data. If a transaction has DB2, CICS, or MQSeries events associated with it, those events are also displayed.

  Detail trace can be used to debug complex application performance problems associated with the execution of a specific transaction. This data can be saved to data sets for later reference or use on another system.

When you trace a transaction, you can:

- Qualify a trace by workload selection or by exception filters to limit the trace; for example, by database name or high elapsed time.
- Set automatic start and stop times for the trace when you think the most activity will occur.
- Use integrated trace display services where you can see a list of transaction trace entries and expand to summarized trace data or detailed trace events for a selected transaction.

Chapter 27. Requesting Workload Trace Data Collection (MTRAC)

An IMS workload trace must be activated by using the MTRAC service. The data collected by MTRAC is displayed by trace display services as described in Chapter 29, "Requesting a Workload Trace Data Display (LTRAC, STRAC, and DTRAC)" on page 239.

When a trace is requested, the data entry panel for the MTRAC service is displayed. Options on this panel let you specify:

- Type of trace

  Three trace types are available:

Summary Trace	The summary trace (TYPE=S) provides timestamps and other important workload statistics about the activity generated by one transaction occurrence
Abbreviated Trace	The abbreviated summary trace (TYPE=A) provides the same information as a summary trace except when the abbreviated trace is logged to a data set. The logged trace will contain only LTRAC entries (no STRAC data), which reduces DASD requirements. You can use the current traces option to view STRAC data while the trace data resides in the buffers.
Detail Trace	The detail trace (TYPE=D) provides more detailed chronological event data about the activity that occurred for one transaction. The amount of data collected is limited to a specified buffer size, as described in "Trace Requirements" on page 217. Some transactions may require multiple logical terminal pages to display all of the traced information. Because of the output volume a detail trace might produce, you should use it only when you are analyzing performance problems in specific applications. For example, a valuable use of detail trace is during new application testing to analyze call flow, I/O, and performance.

  The type of trace is specified by the TYPE keyword with the MTRAC request.

- Trace logging to trace log data sets

  Trace logging allows you to record both summary and detail trace data to external VSAM data sets. It is requested with a yes (Y) option for LOGTRAC (log trace) with the MTRAC trace request. When you ask for trace logging, a Trace Log Data Set Options data entry panel is displayed. These trace logging options let you specify:

  – Trace log data set allocation and disposition, archiving, and switching

  – Whether existing data in a log data set can be overwritten

- Trace data display by selection criteria and exception filters

Selection criteria	You can limit a trace to a specific part (or parts) of the IMS workload by: – Program type – Transaction type – Transaction name – Transaction class

– PSB name
– Region job name
– User ID
– Program name
– Region ID
– Terminal ID
– Database name
– DB2 plan name
– Abend code

Exceptions — You can limit a trace to exception conditions by setting thresholds for time, such as transaction elapsed time or input queue time, or counts, such as number of databases used or number of DC calls. A minimum or maximum threshold can be defined. Only those trace records that meet the defined threshold condition are collected.

Note: For MTRAC to collect DBCTL thread data, CICS=YES or CICS=ONLINE must be specified in BBPARM member IMFECP00.

Starting MTRAC

MTRAC data collection can be activated by:

- Selecting the TRACES option from the Primary Option Menu

  Enter a START TRACE (ST) command in the CURRENT TRACES application selected with the TRACES option. When ST is requested, the initial data entry panel displays a default request for a summary trace (S for the TYPE parameter), as shown in “Start Workload Trace Request Panel” on page 205. Changing TYPE to D requests a detail trace.

  Note: As described in *Implementing Security for MAINVIEW Products*, you must be authorized to request a trace.

  You also can access the MTRAC initial data entry panel by selecting the MTRAC service from the list of monitors displayed with the SM primary command.

  – Pressing ENTER activates a request.
  – Pressing END redisplays the list application where you requested MTRAC.

  Multiple trace requests with different selection criteria can be run concurrently. Each MTRAC request can be given an optional name for identification by specifying a 1- to 8-character identifier in the PARM field. One MTRAC request can be activated without an identifier; the request identifier defaults to blanks.

  For example, this MTRAC data collection request for the transaction PART is identified as MYTRAC1:

```
PARM    ==> MYTRAC1
TYPE    ==> DETAIL
TRAN    ==> PART
```

  Note: A detail trace for a BMP or JBP must be requested before the region is started. For MPPs and JMPs, the detail trace request must be active before the program is scheduled in the region. The first traced transaction may not have detail data although it was requested because the detail trace buffer was not obtained at message GU (GET UNIQUE) time.

- Defining MTRAC requests in a member of your BBI-SS PAS BBPARM data set that can be started automatically when the system starts or at your request (see "Request Initiation" on page 301)

- Starting MTRAC from a MAINVIEW AutoOPERATOR EXEC

  Write an EXEC that starts MTRAC (MAINVIEW AutoOPERATOR must be installed). Use the IMFEXEC IMF command followed by the service name, optional parameters, and an identifier for the target system; for example:

  `IMFEXEC IMFC SET REQ=MTRAC TRAN=PAY+ TARGET=IMSVS`*x*

  For more information, see the *MAINVIEW AutoOPERATOR Advanced Automation Guide for CLIST EXECS*.

- Starting MTRAC from a service display

  Define a SET request for MTRAC in the SERV field as described in Chapter 37, "SET Timer Request" on page 301.

Using the Workload Trace Data Entry Panel

The Workload Analyzer Trace (MTRAC) monitor is a monitoring service that activates either a summary IMS workload trace or a detail transaction processing event trace. There are two ways to display the request data entry panel shown in Figure 15:

- ST command for START TRACES

- S line command for MTRAC from the monitors service list (SM for START MONITORS)

You can use the request data entry panel to access additional data entry panels to override default settings and request:

Trace logging options to record traces to external VSAM data sets

- Selection criteria to limit the amount of data saved by a specific part of the IMS workload

- Exception filters to limit the amount of data saved by exceptions These options are described in this section.

When a trace is requested, the following data entry panel is shown:

```
BMC SOFTWARE ------------- START IMS TRACE REQUEST ----------- PERFORMANCE MGMT
COMMAND ===>                                                 TGT ===> IMST
                                                              TIME --  15:37:49
PARM     ==>          (Trace identifier)     START ==>          (hh:mm:ss)
TYPE     ==> S        (S-Summary,D-Detail    STOP  ==>          (hh:mm:ss/mins)
                       A-Abbreviated Summary)
STORAGE  ==> 256K     (Display buffer size)  WRAP  ==> YES  (Y/N wrap buffer)
LOGTRAC  ==> N        (Y/N log trace)        RST   ==> HOT  (HOT,COLD,PUR,QIS)
TITLE    ==> WORKLOAD TRACE                  QIS   ==> YES  (YES,NO)

Specify additional trace options:     (*=processed)
   Selection criteria         ==> N   (Y/N)
   Exception filters          ==> N   (Y/N)
   Trace log data set options ==> N   (Y/N)

                                       Press ENTER to process; END to cancel
```

Figure 15. Start Workload Trace Request Panel

The options in this panel are SET keywords for an MTRAC request. These keywords are described in Table 8 in alphabetical order.

Table 8. SET Keywords to Activate a Trace

Keyword	Operand	Description
LOGTRAC	N|Y	N is the default. Y writes all trace records for this request to a unique trace log data set (TLDS) for this trace. If Y is specified with no additional summary or detail trace logging options (see Chapter 28, “Logging a Trace” on page 221), a single data set is allocated dynamically using the defaults specified in IMFBEX00 for this MVIMS. **Note:** The BBI-SS PAS started task must be authorized (see the security description in *Implementing Security for MAINVIEW Products*) to allocate trace log data sets dynamically.
PARM	id	Specifies which set of trace data is to be displayed with the application trace services. The id is a unique 1- to 8-character identifier. If you want to run multiple trace requests, specify an ID in the PARM field to make the trace unique. Blank can also be used as an identifier for one trace.
QIS		Defines the action to be taken for the service when IMS is not active.
QIS	YES	Specifies that the service is to be quiesced. This is the default.
QIS	NO	Specifies that the service is to start or continue running. **Note:** When `QIS=NO` is specified, monitors that require IMS continue to be scheduled at each interval; however, they return zero values. Any analyzer services set up to run asynchronously fail with a short message of CANNOT LOCATE IMS SPECIFIED in the first line. The BBI-SS PAS Image log contains screen images of these services.
RST		Defines the restart option to be used restart service when a service is quiesced because of an inactive IMS subsystem or `RRR` is specified for BLK (see “Multiple Requests” on page 302). Default is HOT.
RST	HOT	Restarts the trace automatically without loss of previous data.
RST	COLD	Restarts the trace automatically; all previously collected data is deleted.
RST	PUR	Purges the trace automatically when the target IMS starts.
RST	QIS	Keeps the trace in a quiesced state until it is purged by an authorized user.

Table 8. SET Keywords to Activate a Trace (continued)

Keyword	Operand	Description
START	hh:mm:ss	Requests trace start time. If the time entered is more than 10 minutes prior to the current time, 24 hours are added to the specified time and the request is started the next day. To start a request at midnight, specify 24:00:00. The default is the next full minute. **Note:** This option cannot be modified. The request must be purged and a new request must be made.
STOP	`nn|hh:mm:ss`	Requests trace stop limit, where `nn` is length of time in minutes and `hh:mm:ss` is a timestamp. Processing stops at the end of the last interval before the specified stop time. This time is displayed in the STOP field when the request is viewed with the R, P, M, or W line command from the Current Traces (Primary Menu Option 3) application. If the time entered is the same as the START time, 24 hours are added to the STOP time.
STORAGE|ST	`nnnnnn`	Specifies the amount of BBI-SS PAS private storage to be used by the MTRAC service. MTRAC saves captured data in main storage. Values from 1 to 999999 can be specified, either by byte or by kilobyte; for example, both 40000 and 40K are valid entries. The 40000 byte specification will be rounded to the next multiple of 1024 (1K), in this case 40960 (40K). The MTRAC data entry panel for this option is primed from the default specified in BBPARM member IMFBEX00. If no value is specified, the value from BBPARM BBIISP00 is used. **Note:** For a detail trace, this value must be at least four times TRSIZE specified in BBPARM member BBIISP00 (see “Trace Requirements” on page 217). If STORAGE is less than four times TRSIZE, the STORAGE value is adjusted upward.
TITLE|T	'c...c'	Defines the service display title (1 to 24 characters). A user-defined title can replace the default service title. If the title is specified with SET as a single request or in a BBPARM member as a series of requests, it must be enclosed in single quotation marks. **Note:** The MOD keyword cannot be used to change this option. The request must be purged and a new request must be made.

Table 8. SET Keywords to Activate a Trace (continued)

Keyword	Operand	Description
TYPE	S|SUMMARY D|DETAIL A|ABBREVIATED	Trace event data is in either summarized (SUMMARY, the default) or detailed (DETAIL) format. DETAIL collects information on every DLI, DB2, or MQSeries call. The abbreviated summary trace (TYPE=A) provides the same information as a summary trace except when the abbreviated trace is logged to a dataset. The logged trace will contain only LTRAC entries (no STRAC data). You can use the current traces option to view STRAC data while the trace data resides in the buffers. **Note:** You must be authorized to start a summary or detail trace.

WRAP	YES|
	NO
		Determines trace data wrap in BBI-SS PAS buffer. Default is YES; new data overlays oldest data. NO stops the trace when the buffer is full. A SET request with `MOD=MTRAC, WRAP=YES` resumes the trace. This can be specified in a BBPARM member or as:
		
		`SERV ==> SET`
		`PARM ==> MOD=MTRAC, WRAP=YES`

Filtering a Trace

You can request additional data entry panels to specify:

- Trace log data set options
- Selection criteria
- Trace exceptions by specifying Y for one or more of these additional trace options.

When you select Y for an option, additional data entry panels are displayed. Once you are finished with a panel, press ENTER to process those request options. Each panel you have used is displayed before the Start IMS Trace Request panel is returned.

When all panels are processed, the Start IMS Trace Request panel is redisplayed with an * character next to each processed option. The request is not activated until all the options display either N or *. This gives you the opportunity to review and change your options before sending the request.

For example:

- To eliminate previously specified options, change the * to a Y and clear the options you do not want from the associated data entry panel.
- If you change an * to a Y or an N to a Y, the data entry panel for that option is primed with the values previously specified. You can leave them as shown or change them.

Note: This allows you to change any option you have requested. However, once a trace is activated, you can change only some, not all, of the options as described in “Modify Trace Options Request (M Line Command)” on page 245. Any others must be changed by stopping the trace and reactivating it.

Filtering a Trace by Selection Criteria

Selection criteria limit or filter trace data display to a specific part of the IMS workload. By specifying a certain resource to be traced, such as an IMS region, the display is limited automatically to those events occurring only with that resource.

To limit data display by selection criteria, specify Y in the Selection Criteria field of the Start IMS Trace Request data entry panel. This displays the Selection Criteria data entry panel shown in Figure 16.

```
BMC SOFTWARE  ---------- IMS TRACE SELECTION CRITERIA -------- PERFORMANCE MGMT
COMMAND ===>                                                     TGT ===> IMSA

 Specify selection criteria:
   PGMTYPE  ==>                 (MPP, MDP, IFP, FPU, TPI, BMP, DBT, NOTDBT)
   TRANTYPE ==>                 (ALL, DLI, DB2, FP)
   TRAN     ==>
   CLASS    ==>
   PSB      ==>
   REGION   ==>
   USERID   ==>
   PROG     ==>
   RGNID    ==>
   TERM     ==>                 (name, SYNCLOCK, MSCCLOCK)
   DATABASE ==>
   DB2PLAN  ==>
   ABEND    ==>
```

Figure 16. Selection Criteria Data Entry Panel. Start IMS Trace Request Options

A trace can be requested by IMS workload components by using the workload selection keywords shown in Table 9 on page 210. If none of the selection keywords are specified with an MTRAC request, all transactions that complete processing are traced.

Following is an example of using parameters to selectively trace transactions. The conditions are ANDed.

```
PARM   ==> TRACE1
TRAN   ==> PART
PSB    ==> DFS+++
```

This request traces all transactions with a transaction code of PART and a PSB whose name is six bytes long and begins with the letters DFS.

Note: Conversational transactions are accepted for either `PGMTYPE=MPP` or `PGMTYPE=BMP`.

Syntax

- A list of operands can be specified in each of the selection criteria fields. Each entry in the list must be separated by a comma or blank. For example:

  `TRAN ==> TRAN1,TRAN2,TRAN3`

  traces all transactions that have a transaction code of TRAN1 or TRAN2 or TRAN3.

- The + qualifier character can be used generically or positionally. As a generic qualifier, it cannot be followed by any other character. As a positional qualifier, it must be repeated for every character to be replaced. For example:

  ```
  REGION ==> CICS+
  ```

  is a data collection request for workload trace to sample only regions whose names begin with CICS.

  Entering:

  ```
  TRAN ==> A++N+
  ```

  traces all events that have an identifiable transaction code with an A character in the first position, any two characters before N, and any characters following N.

- Trace records are created when there is a match with the selection criteria keywords. Multiple keywords are connected with an AND operation and the operands within a keyword are connected with an OR operation. For example:

  ```
  PROG    ==>  SAMPPGM,TESTPGM
  PGMTYPE ==>  3,4
  ```

  collects only records that use programs SAMPPGM or TESTPGM and are running in classes 3 or 4.

The options in the IMS Trace Selection Criteria data entry panel are SET keywords for an MTRAC request. These keywords are described in Table 9 in alphabetical order.

Table 9. SET Keywords to Define Trace Selection Criteria

Keyword	Operand	Description
ABEND	+| Sxxx| Uxxxx	Qualifies a trace request by abend code, where xxx is an alphanumeric code and: + Indicates all abends S Indicates a system abend; for example, S0C7 You must specify four characters. U Indicates a user abend; for example, U3303 You must specify five characters. The plus qualifier (+) can be used.
CLASS	class	Qualifies a trace request by the numeric scheduling class(es). Multiple classes can be specified as follows: `CLASS ==> 1 2 3` Default is to include all IMS scheduling classes. **Note:** This workload selection keyword cannot be used for an MTRAC request for DBCTL threads.
DATABASE	database name	Qualifies a trace request by the specified database.
DB2PLAN	DB2 plan name	Qualifies a trace request by the specified plan.

Table 9. SET Keywords to Define Trace Selection Criteria (continued)

Keyword	Operand	Description
PGMTYPE| PTYP	MPP|MDP IFP|FPU TPI|BMP DBT|NOTDBT	Qualifies a monitor request by program type. Multiple program types can be specified. The types are: MPP Message processing program or Java message processing program (JMP) MDP Message-driven program IFP IMS Fast Path program FPU Fast Path utility TPI CPI-C driven program BMP Batch message processing program or Java batch message processing program (JBP) DBT DBCTL threads (CICS and ODBA) Default is all types. DBT selects only programs that use a DBCTL thread. NOTDBT selects programs that do not use a DBCTL thread.
PROG	program name	Qualifies a trace request by the specified program. Valid entries are 1- to 8-character alphanumeric program names, which can be used with the plus qualifier (+). Multiple program names can be specified; for example: `PROG ==> abc cde+`
PSB	psb name	Qualifies a trace request by the specified IMS program specification block. Valid entries are 1- to 8-character alphanumeric PSB names, which can be used with the plus qualifier (+). Multiple PSB names can be specified; for example: `PSB ==> ab1 c+2`
REGION	job name	Qualifies a trace request by the specified job name of an IMS region. Valid entries are 1- to 8-character alphanumeric region job names, which can be used with the plus qualifier (+). Multiple region job names can be specified; for example: `REGION ==> +ba ace`
RGNID	region number	Qualifies a trace request by the specified region ID. Valid entries are 1- to 3-character numeric region IDs from 1 to `nnn`, where `nnn` is a valid region number. Multiple region IDs can be specified; for example: `RGNID ==> 1 23 200`

Table 9. SET Keywords to Define Trace Selection Criteria (continued)

Keyword	Operand	Description
TERM	lterm	Qualifies a trace request by the name of the IMS LTERM. Valid entries are 1- to 8-character alphanumeric LTERM names, which can be used with the plus qualifier (+). Multiple LTERM names can be specified; for example: `TERM ==> L071 L0+`
	SYNCLOCK	Defines a trace request to include *all* transactions from systems synchronized to the clock of the local IMS (in other words, all local transactions, all local MSC transactions, and all nonlocal MSC transactions that arrive from a system with its clock synchronized to the clock of the local IMS).
	MSCCLOCK	Defines a trace request to include *only nonlocal* transactions from MSC systems synchronized to the clock of the local IMS.
TRAN	trancode	Qualifies a trace request by the specified IMS or CICS transaction. Valid entries are 1- to 8-character alphanumeric transaction codes, which can be used with the plus qualifier (+). Multiple transaction codes can be specified; for example: `TRAN ==> TR+ PAY1` For a CICS DBCTL transaction, the CICS transaction code is used. For a Fast Path transaction, the routing code assigned by the DBFHAGU0 exit routine is used.
TRANTYPE|TTYP	ALL|DB2|DLI|FP	Qualifies a trace request by transaction type: ALL All transaction types DB2 Those that access DB2 DLI Those that make DL/I database calls FP Those that make Fast Path database calls
USERID	userid	Qualifies a trace request by the name of the IMS sign-on user ID. Valid entries are 1- to 8-character alphanumeric IMS sign-on user IDs, which can be used with the plus qualifier (+). Multiple user IDs can be specified; for example: `USERID ==> USER1 USER2 USER6`

Filtering a Trace by Exceptions

You can use exception keywords to specify thresholds that determine which trace records are to be saved. These keywords can be used to:

- Detect occurrences of specific exception conditions.
- Limit the amount of trace storage either for online viewing or logging to VSAM data sets by tracing only the most interesting transactions for analysis. To limit data display by exceptions, specify Y in the Exception Filters field of the Start IMS Trace Request data entry panel. This displays the Exception Filters data entry panel shown in Figure 17.

```
BMC SOFTWARE  ----------- IMS TRACE EXCEPTION FILTERS -------- PERFORMANCE MGMT
COMMAND ===>                                                     TGT ===> IMSA

TIME:                      DLI CALLS:               SQL CALLS:
 Elapsed    ==>             Total    ==>             Total      ==>
 Input Q    ==>
 Response   ==>             DB calls ==>             Sel/Fch    ==>
 Tot CPU    ==>             Get type ==>             Open       ==>
 Dep CPU    ==>             Upd type ==>             Insert     ==>
 DB2 CPU    ==>             REPL     ==>             Delete     ==>
                            ISRT     ==>             Update     ==>
 MISC COUNTS:               DLET     ==>             DDL        ==>
 DBs used   ==>                                      Dynamic    ==>
 Read/write ==>             DC calls ==>             Others     ==>
 Serial I/O ==>             Msg GN   ==>
                            Msg ISRT ==>             FP:
                            Msg PRG  ==>              CI contend ==>
                                                      FPwait Bfr ==>
                                                      OBA used   ==>
Transactions that meet one or more of the above thresholds are accepted  for the
trace.  If no filters are specified, all trace records that meet  the selection
criteria are accepted.
```

Figure 17. IMS Trace Exception Filters Data Entry Panel

A trace can be requested by exceptions by using the keywords shown in Table 10 on page 214. If none of the selection keywords are specified with an MTRAC request, all transactions that complete processing are traced.

Syntax

- Exception filter keyword values are user-defined.
- A < prefix can be used with a numeric value to define a less-than condition.

  If you do not specify a less-than condition, a comparison is made to an equal-to or greater-than condition.

  Note: If you want to trace transactions with a 0 exception, you must specify <1 for the exception filter. For example, to check transactions that have made 0 GU calls, specify <1 for the Get Type exception filter. A 0 value specified for any of these options disables that option; the specified exception is not traced.

- Specified keywords are ORed; a trace record is saved if any one of the comparison conditions is valid.

For example:

```
Response  ==> 5.5
```

collects only records with a response time equal to or greater than 5.5 seconds.

```
DB calls  ==> <5
```

collects only records that have fewer than 5 file database calls.

Note: The options in the IMS Trace Exception Filters data entry panel are described in Table 10 on page 214 in alphabetical order and are followed by a corresponding SET keyword. These keywords can be used for a SET request for a service or for multiple SET requests in a BBPARM member, as described in Chapter 37, “SET Timer Request” on page 301.

Time can be specified either in hours, minutes, and seconds as `hh:mm:ss` or in milliseconds as `n.nnn` (8-digit maximum); for example, 3.456 or 3,456 milliseconds or .001 or 1 millisecond. A numeric (`n`) can be specified from 1 to 999999 or 99999K; for example, 100K (multiple of 1024) I/Os.

Table 10. Exception Filter Keyword Descriptions

Option / Keyword	**Operand**	**Description**
CI Contend / FPCICON	`n`	Number of Control Interval contentions
DB2 CPU / DB2CPU	`hh:mm:ss`| `n.nnn`	DB2 CPU time
DB calls / DLICALL	`n`	Number of database calls
DBs used / DBUSED	`n`	Number of databases used
DC calls / DLIDC	`n`	Number of DC (message) calls (message GET UNIQUE, message GET NEXT, for example)
DDL / SQLDDL	`n`	Total number of DDL and SQL calls
Delete / SQLDEL	`n`	Total number of DELETE SQL calls
Dep CPU / DEPCPU	`hh:mm:ss`| `n.nnn`	Dependent region CPU time
DLET / DLIDLET	`n`	Number of DELETE DL/I calls
Dynamic / SQLDYN	`n`	Total number of dynamic SQL calls
Elapsed / ELAPSED	`hh:mm:ss`| `n.nnn`	Transaction elapsed time
FPwait Bfr / FPWTBFR	`n`	Number of Fast Path buffer waits
Get type / DLIGTYP	`n`	Number of GET DL/I calls (GET UNIQUE, GET HOLD UNIQUE, GET NEXT, GET HOLD NEXT, for example)
Input Q / INPUTQ	`hh:mm:ss`| `n.nnn`	Time on the input queue

Table 10. Exception Filter Keyword Descriptions (continued)

Option / Keyword	Operand	Description
Insert / SQLINS	n	Total number of INSERT SQL calls
ISRT / DLIISRT	n	Number of INSERT DL/I calls
Msg GN / DLIMSGN	n	Number of message queue GET NEXT calls
Msg ISRT / DLIMSIN	n	Number of message queue INSERT calls
Msg PRG / DLIMSPR	n	Number of message queue PURGE calls
OBA used / FPOBA	n	Number of overflow buffers used
Open / SQLOPN	n	Total number of OPEN SQL calls
Others / SQLOTH	n	Total number of other SQL calls
Read/write / RDWRT	n	Number of READs and WRITEs
REPL / DLIREPL	n	Number of REPLACE DL/I calls
Response / RESPONS	hh:mm:ss| n.nnn	Response time (elapsed + input queue time)
Sel/Fch / SQLSEL	n	Total number of SELECT and FETCH SQL calls
Serial I/O / SERIO	n	Number of serialized I/Os
Tot CPU / TOTCPU	hh:mm:ss| n.nnn	Total CPU time
Total / DLITOT	n	Total number of DL/I calls
Total / SQLTOT	n	Total number of SQL calls
Upd type / DLIUTYP	n	Number of update DL/I calls (REPLACE, INSERT, DELETE)
Update / SQLUPD	n	Total number of UPDATE SQL calls

Stopping MTRAC

Note: A stopped trace is marked complete and stops collecting trace data, but the data remains available for viewing with the DTRAC service until the MTRAC request is purged. A purge stops data collection and deletes any accumulated data.

An MTRAC request can be stopped by:

- Stopping the request with a Z line command

  Use the Z line command from the Active Timer list application as described in Chapter 6, “Displaying a List of Active Timer Requests (Primary Menu Option 2)” on page 65.

- Setting a stop time for automatic completion of data collection

  Specify the STOP value (as a timestamp) on the start or modify MTRAC request panel or with a SET request for MTRAC from the service display panel. The collected data remains available for viewing until the request is purged. For example, the following SET request:

  ```
  SET
  REQ=MTRAC,MYTRAC1,START=11:00:00,STOP=11:10:00
  ```

  stops MTRAC data collection at the end of 10 minutes. The data remains available for viewing.

- Purging a request from the Active Timer list with the P line command

  Use the P line command from the Active Timer list application as described in Chapter 6, “Displaying a List of Active Timer Requests (Primary Menu Option 2)” on page 65.

- Purging a service request with a SET request

  Issue a PRG request with SET from the service display panel, BBPARM, or a MAINVIEW AutoOPERATOR EXEC (MAINVIEW AutoOPERATOR must be installed); for example:

  ```
  SET
  PRG=reqid|ALL
  ```

The SET parameters are described in Chapter 37, “SET Timer Request” on page 301.

Trace Requirements

A detail trace stores trace data for active transactions in ECSA buffers. The number and size of the detail trace buffers can be specified in either IMFBEX00 or the BBIISP00 member of the BBPARM data set. IMFBEX00 has priority over BBIISP00. The defaults for BBIISP00 are:

```
TRBUFF=10
TRSIZE=4
```

TRBUFF defines the number of detail trace buffers. TRSIZE defines the size, in K, of each trace buffer. For example, if you specify TRSIZE=50000, each trace buffer will be 48K.

When an active transactions completes, the contents of its buffer are moved to the trace areas in the BBI-SS PAS and can then be displayed by using the DTRAC service.

A pool of buffers is GETMAINed in ECSA when a detail trace is activated and then dynamically allocated to the dependent regions as needed. The buffers are returned to the pool when the trace areas in the BBI-SS PAS are updated. If, during transaction initialization, one of the buffers cannot be obtained, only a summary record is generated for the transaction.
The buffer pool is FREEMAINed either when the last detail trace stops (and the program running in each region which has a detail trace buffer allocated terminates) or when the target system terminates. The buffers in ECSA are obtained only if a detail trace is activated.

Note: To calculate an appropriate TRSIZE, estimate about 16K per 100 DL/I or SQL calls. The recommended value for TRBUFF is the maximum number of IMS dependent regions running concurrently plus 2. If there are not enough detail trace buffers, only a summary trace record is created.

MTRAC stores trace data in the private storage area of the BBI-SS PAS address space in extended private storage. This should be considered when setting up storage requirements for the BBI-SS PAS (see the *MAINVIEW Common Customization Guide* for a description of the storage requirements). When the available allocated area is full, the newest data wraps around and overlays the oldest data unless `Wrap=NO` is specified with the MTRAC request.

The size of the trace data storage area in the BBI-SS PAS can be specified with the STORAGE parameter in the MTRAC request data entry panel; for example:

```
STORAGE ==> 100K
```

requests a GETMAIN storage of 100K. The default can be specified in either the IMFBEX00 or BBIISP00 member of the BBI-SS PAS BBPARM data set. IMFBEX00 has priority over BBIISP00.

A summary trace entry requires a minimum of 376 bytes of storage per transaction. The largest amount of storage that can be used is 2776 bytes:

```
376 + (48 * 50) = 2776
```

This is 376 minimum bytes plus 48 bytes (maximum of 50) for each database trailer appended to the summary record (see the DBTS parameter in “Specifying Event Collector Options” in the *MAINVIEW for IMS Online – Customization Guide*.

A detail trace generally requires a much larger storage area than a summary trace; 52 bytes are required per detail line. So a detail trace of a transaction with 20 detail lines requires about `20 * 52 + summary trace` bytes of storage. Specifying additional storage prevents frequent wrapping.

Qualifying a Workload MTRAC Request

To limit the IMS workload components traced but not the resources, a generic operand of + can be specified for an MTRAC keyword to select all resources of that type. For example, `REGION=+` traces all regions but excludes input and output events. `TRAN=+` traces all events that have an identifiable transaction code associated with them, which eliminates output queuing and output communications.

The + character can be used generically or positionally:

- As a generic resource name qualifier, it cannot be followed by any other character.
- As a positional qualifier, it must be repeated for every character to be replaced.

For example:

```
REGION ==> CICS+
```

is a data collection request for workload trace to display only those regions whose names begin with CICS.

Entering:

```
TRAN ==> A++N+
```

traces all events that have an identifiable transaction code with an A character in the first position, any two characters before N, and any character following N.

The + qualifier can be used for MTRAC service requests for:

Resource Selection	Resource
PROG keyword	Program names
PSB keyword	PSBs
REGION keyword	Region job names
TERM keyword	LTERM names
TRAN keyword	Transaction names
USERID keyword	User IDs
DATABASE keyword	Database

Note: The generic qualifier is not valid for the CLASS, PGMTYPE, or RGNID keywords.

Controlling Active Traces with CURRENT TRACES (Primary Menu Option 3)

You can control active trace requests by selecting the Current Traces application (Primary Menu Option 3) from the Primary Option Menu as described in Chapter 30, “Displaying a List of Current Traces (Primary Menu Option 3)” on page 241. From a list of current traces displayed by this application, you can use line commands to:

- Display collected trace data
- Show trace options
- Modify a trace request
- Switch trace log data sets
- Purge a trace request
- Replicate a trace request
- Quiesce trace logging

You can also enter primary commands to start a trace (ST) or display the History Traces (HT) application which is used to display trace log data set requests described in Chapter 28, “Logging a Trace” on page 221. For information about how to use the Current Traces application, see Chapter 30, “Displaying a List of Current Traces (Primary Menu Option 3)” on page 241.

Displaying an Online Trace with CURRENT TRACES (Primary Menu Option 3)

You can use the S line command from the Current Traces application to select an online trace for display. The LTRAC display for the selected trace is shown. You can use the LTRAC display to select an entry to see summarized trace data for a transaction instance (STRAC display). From here, you can select the DTRAC display to see detail trace data for that transaction instance. For information about how to use these displays, see Chapter 29, “Requesting a Workload Trace Data Display (LTRAC, STRAC, and DTRAC)” on page 239

Chapter 28. Logging a Trace

When you request a trace with the MTRAC service, you can also ask for the trace data to be recorded to external VSAM data sets called trace log data sets (TLDS) for later viewing. These external VSAM data sets can store both summary and detail traces. Trace log data sets can be dynamically allocated by user-selected trace log data set options with the MTRAC request or they can be pre-allocated using sample batch jobs. They are managed from the HISTORY TRACES application.

The logging of trace data requires a pre-allocated trace directory that must be identified to BBI. Setup of the directory can be done automatically during AutoCustomization or manually. The online trace directory keeps track of all trace log data sets and is maintained through the HISTORY TRACES application selected from the Primary Option Menu. The HISTORY TRACES application also can be used to display any trace logged to the trace log data set(s).

The MTRAC trace log data set options are used to assign a set of logs with each MTRAC request. These options also allow:

- Log switching
- Data set archiving when a data set is full or it switches to another data set or a trace completes normally
- Data overwrite of a data set that has not been reset yet

Log switching to the next available data set can be requested to occur:

- When a log data set is full
- At a specific time of day
- Immediately with a line command from the CURRENT TRACES application

Logging switches to the oldest trace data set if a data set is not available and the MTRAC request specifies overwrite.

If archiving is requested for a trace data set, an archive started task is run. This started task:

- Makes a copy of the data to tape or another data set
- Executes a step that allows the data set to be reused

If your site prefers to use batch jobs to pre-allocate and maintain the trace directory and the trace log data sets manually instead of dynamically, there are sample jobs available in the BBSAMP library, as described in “Maintaining Trace Data Sets Using Sample Batch Jobs” on page 229. Pre-allocated trace log data sets must have a low-level qualifier of V*nn*, where *nn* can be 01 through 99. The first data set must have a low-level qualifier of 01.

Using Trace Logs

Trace logging provides:

- Log data sets for each trace request

  Each trace request can be assigned its own log data sets. This allows dynamic application traces to be managed separately from system-wide, continuous workload surveillance or system exception traces.

- Flexible space management

  Each trace can be logged to single or multiple data sets either pre-allocated manually or dynamically allocated and managed by the online trace services.

- User-selected logging options

  A trace can be logged simply by using option defaults, or all the options can be user-defined for complete control.

- Viewing of active trace data online

  The trace displays (LTRAC, STRAC, DTRAC) can be used to see the current data from an active trace online or from one being logged to a data set.

- Viewing of historical trace data at any time

  The History Traces application (Primary Menu Option 4) displays the data sets tracked in the online trace directory. It allows online viewing of the data recorded in these data sets with a simple line command whether the trace was written the same day, last night, a year ago, or even on another system. This means:

  – Data from completed traces saved in the trace logs can be viewed later; there is no lost data because of system interruptions.

  – Older data from an active trace with multiple log data sets can be accessed.

  – Traces can be kept as transaction profiles for later comparisons if application or system changes affect system performance.

- Tracing of very large applications

  Long-running applications can be traced without loss of data because of online trace storage buffer limitations.

Defining a Trace Directory

Before any trace can be started with trace logging, a trace directory must be:

- Allocated and initialized

  This can be done during AutoCustomization or manually using BBSAMP member JXT001. The instructions for using JXT001 are in the *MAINVIEW for IMS Online – Customization Guide* and are commented in JXT001.

- Identified to BBI

  It is identified to BBI by the TRDIR parameter in BBPARM member BBIISP00.

There is one trace directory per BBI-SS PAS (identified to that BBI-SS PAS in BBPARM member BBIISP00 as described above) that is dynamically allocated when the BBI-SS PAS starts. It is a VSAM linear data set containing one entry for each trace log data set. Each entry indicates the date and time of data set creation, the current status of the data set, the trace target, and other related information. Entries can be added to or deleted from the directory by using the History Traces application (see Chapter 31, “Displaying History Traces (Primary Menu Option 4)” on page 247).

Requesting Trace Logging with MTRAC Options

You can request trace logging by:

- Entering the ST application transfer command in any display to access the MTRAC request data entry panels.

or

- Selecting MTRAC from the list of monitors displayed with the SM application transfer command.

then

- Entering Y for the LOGTRAC option in the MTRAC data entry panel.

The Trace Log Data Set Options data entry panel, shown in Figure 18 on page 224, is displayed. The options are described in Table 11 on page 225 in alphabetical order. When you are finished specifying the options you want, press the ENTER key to process them. If you have requested any other MTRAC options, each panel you have used is displayed before the Start IMS Trace Request panel is returned. This gives you the opportunity to change any options you previously specified before the trace is activated.

These options will allocate the trace log data sets dynamically, overwrite existing data, archive a log data set when it is full, and switch to the next data set if they are requested. If you want to use batch jobs to manually define, archive, and restore trace log data sets, see “Maintaining Trace Data Sets Using Sample Batch Jobs” on page 229.

```
BMC SOFTWARE ------------ TRACE LOG DATA SET OPTIONS --------- PERFORMANCE MGMT
COMMAND ===>                                                TGT ===> IMSA
                                                             TIME --  12:43:31

Number of logs   ===> 1
First log DSN    ===>
                Low level qualifier of DSN must be V01
                Blank for default: CIR7.IMSA.BLANK.mmmdd.Thhmm.V01
                Names without quotes will be prefixed with MVIQA

Overwrite logs   ===> Y      (Y/N) (Action when all logs used)
Archive PROC     ===>              (Blank for none/PROCLIB member name)
Log switch time  ===>              (HH:MM that a log switch is requested)

Disposition      ===> NEW    (OLD/NEW) If NEW, specify options below

Volumes          ===> (V00001,V00002,V00003)
Primary CYLS     ===> 5                SMS Storage Class     ===> SMSSTOR
Data DSN suffix  ===> D                SMS Data Class        ===> SMSDATA
                                       SMS Management Class  ===> SMSMGMT
```

Figure 18. Trace Log Data Set Options Data Entry Panel

Table 11 describes each option shown in Figure 18 and lists them in alphabetical order.
A corresponding keyword is shown with each option in the table. These keywords can be used with a SET request, described in Chapter 37, “SET Timer Request” on page 301. You can define defaults for the following data set options in BBPARM member IMFBEX00.
The keywords are shown in parentheses.

Overwrite logs (TRREUSE)
Volumes (TRVOLS)

Primary CYLS (TRCYL)
Data DSN suffix (TRSUFFIX)
SMS Storage Class (TRSMSSCL)
SMS Data Class (TRSMSDCL)
SMS Management Class (TRSMSMCL)

Table 11. Trace Log Data Set Options

Option/ Keyword	**Operand**	**Description**
Archive PROC TRARCSTC	`name`	Name of the entry in PROCLIB to archive a log data set when: • It is full • Logging switches to the next data set • Trace completes normally To set up a started task for trace log data set archiving, you can use the sample job in BBSAMP member IMFTARC as described in “Archiving a Trace Log Data Set” on page 230. To restore an archived trace for online viewing, use the sample job described in “Restoring an Archived Trace Log Data Set” on page 230.
Data DSN Suffix TRSUFFIX	`xxx`	Where `xxx` is the suffix to be appended to the name of the VSAM cluster to make the name of the data component. The default value is specified with the TRSUFFIX keyword in BBPARM member IMFBEX00. The initial default is D.
Disposition TRDISP	<u>NEW</u>|OLD	NEW There are no existing trace log data sets. All data sets are allocated when the trace request is processed. If any data set cannot be allocated, the trace request fails and any data set successfully allocated is deleted. OLD There are existing trace log data sets. For example, there may be some data sets that were pre-allocated with BBSAMP member JXT011 as described in “Maintaining Trace Data Sets Using Sample Batch Jobs” on page 229. Any existing data in these data sets is to be overwritten. **Note:** If you are not authorized for dynamic trace log data set allocation (NO is specified for TRALLOC in your BBPARM user authorization member), only OLD is accepted.

Table 11. Trace Log Data Set Options (continued)

Option/ Keyword	Operand	Description
First Log DSN TRDSN	name.V01	Name of the first log data set to be allocated. The low-level qualifier must be V01. If this keyword is omitted and the data set disposition is specified as NEW, a default name is generated as: trprefix.imsid.traceid.mmmdd.Thhmm.V01 trprefix — Your user ID or the prefix specified with the TRPREFIX parameter in IMFBEX00. imsid — ID of the IMS target. traceid — The request identifier specified in the PARM field of the first Start IMS Trace Request panel. If the PARM field is blank, BLANK is the default value for traceid. mmmdd — Current month and day. Thhmm — Current time. V01 — Required suffix. If the name is enclosed in single quotes, it is used as specified. If quotes are not used, the TRPREFIX value specified in IMFBEX00 is added in front of the specified name. If TRPREFIX was not specified, the ID of the user requesting the trace is used. For example, you can specify: TEST.V01 to allocate userid.TEST.V01, where userid is your user ID. If you are using a SET request for MTRAC, as described in Chapter 37, “SET Timer Request” on page 301, instead of the MTRAC request data entry panels and you specify TRDSN without LOGTRAC, Y is assumed for LOGTRAC.
Log switch time TRSWTIME	hh:mm	Where hh:mm specifies the time an automatic log switch from the current log to the next log will occur. This is valid only if the Number of Logs option is >1. **Note:** A switch to another log data set can be made at that time only if there is one available that is empty or marked for reuse.

Table 11. Trace Log Data Set Options (continued)

Option/ Keyword	**Operand**	**Description**
Number of Logs TRNUMDS	n	Where n is the number of trace log data sets to be used for this trace. The default is 1. Specify more than one for continuous traces. This allows automatic switching when a log data set is full. If you are using a SET request for MTRAC, as described in Chapter 37, “SET Timer Request” on page 301, instead of the MTRAC request data entry panels and you specify TRNUMDS without LOGTRAC, Y is assumed for LOGTRAC.
Overwrite logs TRREUSE	Y|N	Indicates the action to be taken when a log data set is about to be used for trace data and is not reset: Y Overwrite existing data N Quiesce logging The default value is specified with the TRREUSE keyword in BBPARM member IMFBEX00. The initial default is Y. **Note:** If the data set is reset, this option has no effect.
Primary CYLS TRCYL	n	The variable n can be 1 to 999. It is the primary allocation in cylinders for trace log data sets. The default value is specified with the TRCYL keyword in BBPARM member IMFBEX00. The initial default is 3.
SMS Data Class TRSMSDCL	name	Name of the SMS data class to be used to allocate a trace log data set. The default is specified in IMFBEX00 with the TRSMSDCL keyword.
SMS Management Class TRSMSMCL	name	Name of the SMS management class to be used to allocate a trace log data set. The default is specified in IMFBEX00 with the TRSMSMCL keyword.

Table 11. Trace Log Data Set Options (continued)

Option/ Keyword	Operand	Description
SMS Storage Class TRSMSSCL	`name`	Name of the SMS storage class to be used to allocate a trace log data set. The default is specified in IMFBEX00 with the TRSMSSCL keyword.
Volumes TRVOLS	`(x,x...)`	Where `x` is the volume(s) to use to allocate a trace log data set. You can specify up to seven volumes. Parentheses are required to specify multiple volumes. The initial default value specified in IMFBEX00 with the TRVOLS keyword is SYSDA. If the Number of Logs option is 1, the data set is allocated on the first volume. If more than one log is specified for the Number of Logs option, the data sets are allocated alternately on each specified volume (if there is enough space) until all data sets are allocated. **Note:** A volume specification replaces any SMS class specification. However, if your site is SMS-controlled, the volume specified may be overridden by your SMS criteria. SMS-controlled sites do not need to specify volume(s). Non-SMS-controlled sites must specify volume(s) to avoid an IDCAMS define error.

Maintaining Trace Data Sets Using Sample Batch Jobs

You can create trace log data sets dynamically as described in "Requesting Trace Logging with MTRAC Options" on page 224 or manually by using sample jobs from the BBSAMP data set. These sample jobs allow you to:

- Create and maintain the trace directory

  As described in "Defining a Trace Directory" on page 223, trace logging requires a pre-allocated trace directory. The BBSAMP sample members you can use for the trace directory are:

  JXT001 — You do not need to use this job if the directory was created by AutoCustomization. It defines and initializes the trace directory. The instructions are commented in JXT001 and are also in the *MAINVIEW for IMS Online – Customization Guide*.

  JXT003 — Synchronizes trace directory information with the actual status of the trace log data sets. The instructions for using this job are commented in JXT003 and are also in the *MAINVIEW for IMS Online – Customization Guide*.

- Define a trace log data set

  As described in "Requesting Trace Logging with MTRAC Options" on page 224, you can let the BBI-SS PAS allocate trace log sets for you, if you are authorized for this, or you can use the following sample job to allocate them manually:

  JXT011 — Defines a trace log data set. The instructions for using this job are commented in JXT011 and are provided in this section.

- Archive a trace log data set

  To set up a started task for automatic archiving, as described in "Requesting Trace Logging with MTRAC Options" on page 224, use:

  IMFTARC — Copies a trace log data set to a sequential data set for archival and marks the trace log data set for reuse. The instructions for using this job are provided in this section.

- Restore an archived trace log data set

  IMFTRLOD — Reloads an archived copy of a trace to a new trace log data set so that it can be viewed online. The instructions for using this job are provided in this section.

- Create a trace log data set from the IMS log

  IMFLOGTR — Creates a summary trace from the IMS log and records it in a trace log data set. The instructions for using this job are provided in this section.

Defining a Trace Log Data Set

You can define different trace log data sets as often as you need them with a batch job or you can let the BBI-SS PAS allocate them for you dynamically, if you are authorized, by using the MTRAC service. You must define them manually if you are not authorized for BBI-SS PAS dynamic data set allocation. BBSAMP member JXT011 is a sample batch job you can use to define a trace log data set (TLDS) as a linear data set using IDCAMS. This pre-allocates a TLDS that can then be referenced when you request a trace with MTRAC.

This job has three steps:

1. Complete the job statement information.
2. Define the name of the trace log data set.

   The low-level qualifier must be `Vnn`, where `nn` can be 01 through 99. The first data set must have a low-level qualifier of V01.
3. Provide a valid volume value for the define statement.

Archiving a Trace Log Data Set

Use BBSAMP member IMFTARC to set up a started task for automatic archiving of a TLDS. It is executed when a data set is full, when logging switches to another data set, or when the trace completes by specifying it with the Archive PROC option on the initial MTRAC data entry panel. This procedure has two steps:

1. Unload the trace log data set to a GDG (generation data group) data set.
2. `PGM=JXTRES` marks the trace log data set as RESET in the header.

You also can use this member as a model to create a batch job to archive or reset a log data set instead of using a started task. It can be invoked from the History Traces panel with an A line command if it is specified with the Archive PROC option on the initial MTRAC data entry panel. Otherwise, it must be submitted manually.

Restoring an Archived Trace Log Data Set

Use BBSAMP member IMFTRLOD to restore an archived trace log data set. This job has two steps:

1. Define a VSAM linear data set.
2. Load the archived data to the defined linear data set.

You can add the linear data set to the online trace directory to view the contents online by using the NEW primary command in the History Traces application (see “NEW Primary Command” on page 251).

Creating a Trace Log Data Set from the IMS Log

Use BBSAMP member IMFLOGTR to create a user-selected summary trace from the IMS log. This job has two steps:

1. Define a trace log data set
2. Load user-selected data from the IMS log to the defined trace log data set

You can add the trace log data set to the online trace directory to view the contents online by using the NEW command in the History Traces application (see “NEW Primary Command” on page 251).

IMFLOGTR defines a trace log data set as a linear data set using IDCAMS. It provides trace identification parameters and filter parameters and produces a report called the “Processing and Exception Log”. This report contains statistics showing all the parameters specified and the results of IMS log processing by IMFLOGTR.

Job Control Language Statements

- Using a copy of BBSAMP member IMFLOGTR, change:
  - The JOB statement.
  - PFX in the PROC statement to your site's requirements.
  - The IMSLOG symbolic parameter to specify your existing IMS log data set and the TLDS symbolic parameter to specify your new trace log data set.
  - The DEFINE statement to identify your new trace log data set (NAME), the volume (VOLUMES) where it is to reside, and the amount of space (CYL) it needs.
- Review IMFLOGTR PARM options for the EXEC statement:

```
//LOAD   EXEC PGM=IMFLOGTR,REGION=4M,PARM='option'
```

Where `option` can be:

READALL
: Read the entire IMS log even if the trace log data set is full. The default is to stop reading the log when the data set is full.

ALLCOUNTERS
: Report the results of IMS log processing and all types and number of log records read. The default is to report only Event Collector transaction records (X'FA'). Other IMS log record counts are not reported.

- Specify the following DD statements for:

SYSPRINT
: Defines the print output data set for the “Processing and Exception Log” report.

SYSUT1
: Defines the IMS log data set.

SYSUT2
: Defines the output trace log data set.

SYSIN
: Defines the IMFLOGTR parameter statements.

- Specify IMFLOGTR parameter statements:

The following keywords are used to identify the trace and filter the summary trace data. Exception filters for IMFLOGTR are the same as those used for an MTRAC request.

Trace Identification Parameters:

SSID	1- to 4-character ID of the BBI subsystem for the trace. The default is the first SMF ID found in the IMS log.
TARGET	1- to 8-character job name of the target system for the trace. The default is the first IMS job name found in the IMS log.
UID	1- to 8-character ID of the user creating the trace. The default is the name of the job creating the trace.

TITLE — 1- to 24-character trace title. The default is LOAD FROM IMS LOG. If your title contains blank characters, it must be enclosed in single quotation marks.

PARM — 1- to 8-character identifier assigned to the trace to make this trace unique. The default is BATCHLOD.

Trace Filter Parameters:

The following parameters qualify the summary trace records from the IMS log by start and stop times and exceptions:

Start and Stop Filters

You can specify:

– Date and time as `yyyy-mm-dd-hh.mm.ss`

– Date as `yyyy-mm-dd`

For example, `START=2000-12-05`

START — Date or date and time to start the summary trace load from the IMS log to the trace log data set. The default is to start with the first record found on the IMS log.

STOP — Date or date and time to stop the summary trace load from the IMS log to the trace log data set. The default is to stop at the last record found on the IMS log.

RDSTOP — Date or date and time to stop the read of the IMS log. If RDSTOP is specified, reading stops when RDSTOP is encountered; no more records are loaded into the trace log data set.

To find out what the total number of summary records in the IMS log is, specify RDSTOP to force processing of the log even if the trace log data set is full. You can use this information to help you size your trace log data set.

Exception Filters

These are the same parameters used by MTRAC to filter a trace by exceptions. As described in “Filtering a Trace by Exceptions” on page 213, you can use these parameters to set a threshold by time or by a count. You can specify a threshold with an = (equal-to or greater-than condition) or a < (less-than condition) character as shown in the examples below.

Exceptions by time can be specified as either `hh:mm:ss` or seconds as `nnn.nnn`, for example,

```
INPUTQ=00:15:00
```

Only those records with an input queue time equal to or greater than 15 seconds are loaded from the IMS log into your trace log data set.

```
RESPONS<.001
```

Only those records with a response time that is less than one thousandths of a second are loaded from the IMS log into your trace log data set.

Exceptions by count can be:

```
DBUSED=2
```

Only those records that have two or more used databases are loaded from the IMS log into your trace log data set.

```
RDWRT<5
```

Only those records that have less than five reads and writes are loaded from the IMS log into your trace log data set.

You cannot specify a decimal place with a count filter.

Time Exceptions

Parameters used as filters to qualify the trace by time include:

DB2CPU	DB2 CPU time
DEPCPU	Dependent region CPU time
ELAPSED	Transaction elapsed time
INPUTQ	Time on the input queue
RESPONS	Response time (elapsed + input queue time)
TOTCPU	Total CPU time

Count Exceptions

Parameters used as filters to qualify a trace by amount of activity against a resource include:

DBUSED	Number of databases used
DLICALL	Number of database calls
DLIDC	Number of DC (message) calls (message GET UNIQUE, message GET NEXT, for example)
DLIDLET	Number of DELETE DL/I calls
DLIGTYP	Number of DL/I calls (GET UNIQUE, GET HOLD UNIQUE, GET NEXT, GET HOLD NEXT, for example)
DLIISRT	Number of INSERT DL/I calls
DLIMSGN	Number of message queue GET NEXT calls
DLIMSIN	Number of message queue INSERT calls
DLIMSPR	Number of message queue PURGE calls
DLIREPL	Number of REPLACE DL/I calls
DLITOT	Total number of DL/I calls
DLIUTYP	Number of update DL/I calls (REPLACE, INSERT, DELETE)
FPCICON	Number of Control Interval contentions
FPOBA	Number of overflow buffers used
FPWTBFR	Number of Fast Path buffer waits
RDWRT	Number of READs and WRITEs

SERIO	Number of serialized I/Os
SQLDDL	Total number of DDL and SQL calls
SQLDEL	Total number of DELETE SQL calls
SQLDYN	Total number of dynamic SQL calls
SQLINS	Total number of INSERT SQL calls
SQLOPN	Total number of OPEN SQL calls
SQLOTH	Total number of other SQL calls
SQLSEL	Total number of SELECT and FETCH SQL calls
SQLTOT	Total number of SQL calls
SQLUPD	Total number of UPDATE SQL calls

Messages

Complete descriptions of the messages issued by IMFLOGTR can be seen by selecting the MESSAGES option (Option M in General Services) from the Primary Option Menu. The last character of the message indicates its type as shown below. The highest severity type is listed first.

C Critical, return code of 12

E Exceptional, return code of 8

W Warning, return code of 4

I Informational, return code of 0

Controlling Trace Logging with CURRENT TRACES (Primary Menu Option 3)

You can control trace requests for logging by selecting the Current Traces application (Primary Menu Option 3) from the Primary Option Menu as described in Chapter 30, “Displaying a List of Current Traces (Primary Menu Option 3)” on page 241. This application displays a list of current trace requests. The line commands you can use in this application to control trace logging include:

- Show trace log options
- Modify a trace log request
- Switch trace log data sets
- Quiesce trace logging

The show line command produces additional logging information, which is described in this section. This section also describes considerations for modifying a trace log request. For information about how to use the Current Traces application, see Chapter 30, “Displaying a List of Current Traces (Primary Menu Option 3)” on page 241.

Show a Trace Log Request (W Line Command)

You can use the W line command from the Current Traces application (Primary Menu Option 3) to view the trace log data set options specified for a trace request.

The IMS Trace Request panel is shown for the selected MTRAC request. To see the log options specified, enter a Y for “Trace Log Data Set Options”. This displays the Trace Log Data Set Options and additional trace log fields read from the current trace log data set, as shown in Figure 19.

```
BMC SOFTWARE  ----------- TRACE LOG DAT SET OPTIONS --------- PERFORMANCE MGMT
 COMMAND ===>                                                      TGT -- IMSA

  Number of logs:       3
  Log DSN:              CIR1.IMSA.SUMMARY.MAY21.T1125.V01

  Recording Start:      15DEC00 00:01:00   Percent used: 100   Status: RESET
  Recording Stop:       15DEC00 23:59:59   Data lost:    No

  Overwrite logs:       Y
  Archive PROC:         ARCHJCL1
  Log switch time:      24:00:00

  Disposition:          NEW

  Volumes:              BAB003,BAB303,BAB009
  Primary CYLS:         3                 SMS Storage Class:
  Data DSN Suffix:      DATA              SMS Data Class:
                                          SMS Management Class:
```

Figure 19. Show Trace Log Data Set Options

This panel shows the previously defined trace log options (see Figure 18 on page 224) and additional fields read from the current trace log data set. The Log DSN field is either the name of the current log data set if the trace is active or the name of the last data set used if the trace is complete.

The fields read from the current trace log data set include:

Recording Start:	Date and time this data set was activated for logging.
Recording Stop:	Date and time of the last record written to this data set.
Percent Used:	Number of pages used divided by the number of pages allocated, expressed as a percent, for this data set.
Data Lost:	YES|NO. Indicates whether any data was lost while writing to this log data set.
Status:	The status of the log data set. This can be:

EMPTY	Data set contains no trace records.
UPDATE	Data set is being updated.
USED	Data set contains valid trace data.
RESET	Data set contains valid trace data and is marked for reuse.
INCOMP	Data set did not close or archive successfully after an update; for example, in a system failure.

Modify a Trace Log Request (M Line Command)

As described in Chapter 30, “Displaying a List of Current Traces (Primary Menu Option 3)” on page 241, you can change MTRAC request options, including those for logging a trace, by selecting a trace request with the M line command. This displays a data entry panel(s) with the options that were defined to activate the selected trace (see Chapter 27, “Requesting Workload Trace Data Collection (MTRAC)” on page 203). Only those options with an `===>` can be changed.

Note: You cannot modify the LOGTRAC option after the trace has started. Thus, you cannot start or stop logging in the middle of a trace with a modify line command. To change the log options specified, enter a Y for “Trace Log Data Set Options”.

The following table describes trace log options you might consider changing when logging is active. All the trace log options are described in Table 11 on page 225.

Table 12. Modifiable Trace Logging Options

Option	Keyword	Possible Use
The first column lists options in alphabetical order. The second column contains the corresponding keyword. The keyword can be used for a SET request or may define a log option default for the data entry panel in BBPARM member IMFBEX00.		
Archive PROC	TRARCSTC=`name`	To change the name of the archive job, perhaps to archive to a different set of archive data sets.
Log switch time	TRSWTIME=`hh:mm`	Allows you to change the time when a log switch occurs. For instance, you might want to add a data set by increasing the number of them (TRNUMDS) and then force a switch to free up the current data set.

Table 12. Modifiable Trace Logging Options (continued)

Option	Keyword	Possible Use
Number of logs	TRNUMDS=n	To increase or decrease the number of trace log data sets. If you increase the number of data sets (TRNUMDS) and NEW is specified for the data set disposition (TRDISP=NEW), additional data sets are allocated before the modify request is considered successful.
Overwrite logs	TRREUSE=Y|N	Changing from Y to N prevents overlaying data that is not archived. Changing from N to Y assures that a data set is always available for new trace data, but old trace data may be overlaid before it can be archived.
Primary CYLS or SMS Data Class SMS Management Class SMS Storage Class Volumes	TRCYL=n or TRSMSDCL=name TRSMSMCL=name TRSMSSCL=name TRVOLS=(x,x...)	If NEW is specified for the data set disposition (TRDISP=NEW) and you are adding log data sets by increasing their number (TRNUMDS), you may want to change the allocation parameters for these new data sets to allocate them on different volumes or to increase their size.

Note: If the number of logs is increased using the modify command, the additional log data sets must be successfully allocated before the request is considered successful. If they cannot all be allocated, the request fails and any additional data sets that were allocated dynamically are deleted. Logging continues with the original number of log data sets.

Switch Log Data Sets (I Line Command)

You can request that trace logging be switched immediately to the next available data set by selecting that trace request with the I line command. An archive request for the trace log data set just deactivated is started if the selected MTRAC request specified the archive option.

If the switch command is issued and only one log data set is defined, logging either reuses the same data set or quiesces. Logging resumes with the same data set if the Overwrite Log option for the MTRAC request was Y (yes). Logging quiesces if the Overwrite Log option for the MTRAC request was N (no).

Note: You must be authorized to issue the I (SWITCH) line command for traces other than your own. This is done with the PMACC parameter (refer to *Implementing Security for MAINVIEW Products).*

Quiesce Trace Logging (Q Line Command)

Once quiesced, logging cannot be restarted for this request. The trace continues to run, but no data is written to the log. To restart logging, the request must be purged with the P line command. Another request for this trace with the logging options specified must be activated.

Note: You must be authorized to issue the Q (QUIESCE LOGGING) line command for traces other than your own. This is done with the PMACC parameter (refer to *Implementing Security for MAINVIEW Products).*

Displaying a Logged Trace with History Traces (Primary Menu Option 4)

As described in Chapter 31, “Displaying History Traces (Primary Menu Option 4)” on page 247, you can display a logged trace from a trace log data set by using the S line command to select a trace log data set from the History Traces list application. This brings up the LTRAC display, described in Chapter 34, “LTRAC - List of Trace Entries” on page 265. When the trace data is from a trace log data set, the following message appears in line 4 of the display:

```
ENTRIES IN DATASET
```

The EXPAND line shows HISTORY. You can use this field to toggle between the data from the selected trace log data set and the current trace data in online storage buffers. Move the cursor to the HISTORY field and press ENTER to see current trace data. Press ENTER to see the selected historical trace data. The target for displayed history trace data is shown in the HIST TGT field of the LTRAC field.

You can use the LTRAC display to select an entry to see summarized trace data for a transaction instance (STRAC display). From here, you can then select the DTRAC display to see detail trace data for that transaction instance. For information about how to use these displays, see Chapter 29, “Requesting a Workload Trace Data Display (LTRAC, STRAC, and DTRAC)” on page 239.

Chapter 29. Requesting a Workload Trace Data Display (LTRAC, STRAC, and DTRAC)

Data collected by an MTRAC request, whether it is still in online trace storage buffers or saved in trace log data sets (see Chapter 28, “Logging a Trace” on page 221), can be displayed by:

LTRAC
: Displays a list of transaction instances selected for a trace request. Each instance corresponds to a completed transaction.

STRAC
: Displays summarized trace data about the activity generated by one transaction instance.

DTRAC
: Displays a detail trace of events that occurred for one transaction in chronological order and includes associated database I/O data and segment search argument, key feedback, and I/O area data. If a transaction has DB2, CICS, or MQSeries events associated with it, those events are also displayed.

Data collection must first be activated with an MTRAC request before any trace display service can be requested. If an identifier is specified with the MTRAC request, that same identifier must be used with the trace display service to view that trace request. Only one MTRAC request can be active without an identifier; the request identifier defaults to blanks.

The easiest way to display a trace is to use an S line command to select an active trace request or a logged trace data set from a service application list. This displays the LTRAC service for that selection. An active request can be selected from the Current Traces application (Primary Menu Option 3), as described in Chapter 30, “Displaying a List of Current Traces (Primary Menu Option 3)” on page 241. A logged trace data set can be selected from the History Traces application, (Primary Menu Option 4), as described in Chapter 31, “Displaying History Traces (Primary Menu Option 4)” on page 247. The Active Timer Requests list application also can be used to select an active trace request.

When you display the LTRAC service for a selected MTRAC request or a logged data set, you can use the trace displays to build trace information as follows:

1. Select Primary Menu Option 3 if you want to see a list of active traces; select Primary Menu Option 4 if you want to see a list of trace log data sets.
2. Use the S (SELECT) line command.

   The LTRAC service showing a list of transaction instances is displayed for that selection.
3. Use the LTRAC S line command to select a transaction instance.

   The STRAC service showing summarized activity for that selected transaction instance is displayed.
4. Use the STRAC EXPAND line to select the DTRAC service for that transaction instance.

   The DTRAC service displays a detailed trace of events in chronological order for that transaction instance and includes associated database I/O data and segment search argument, key feedback, and I/O area data. If a transaction has DB2, CICS, or MQSeries events associated with it, those events are also displayed.

The type of trace shown depends upon the type requested with MTRAC (SUMMARY or DETAIL).

Any of the trace display services for an active MTRAC request also can be requested from a list of analyzer display services (Primary Menu Option 1) or from the SERV field of a display. You must specify the identifier of the MTRAC request when you request the trace display service. For example, for the following MTRAC request identified as MTRAC1:

```
PARM     ==> MYTRAC1
TYPE     ==> DETAIL
TRAN     ==> PART
```

DTRAC is requested for this active trace with the same MYTRAC1 identifier from the list of analyzer display services or SERV field of a display as follows:

```
SERV     ==> DTRAC
PARM     ==> MYTRAC1
```

Chapter 30. Displaying a List of Current Traces (Primary Menu Option 3)

The trace services (MTRAC to request a trace and LTRAC, STRAC, and DTRAC to display the collected trace data) are provided by the Workload Analyzer component of MVIMS.
To view a list of traces activated with an MTRAC request:

- Select Primary Menu Option 3, CURRENT TRACES, from the Primary Option Menu.
- Enter a CT command to transfer to this application from another application. A scrollable list of all the active traces that you are authorized to view is displayed as shown below.

```
BMC SOFTWARE  ------------------ CURRENT TRACES -------------  PERFORMANCE MGMT
COMMAND ===>                                                       TGT ===> IMSA
                                   INPUT    INTVL ==> 3      TIME --  14:31:36
 COMMANDS: ST (START APPLICATION TRACE), HT (HISTORY TRACES), TYPE
 LC CMDS:  S (SELECT), W (SHOW),      M (MODIFY), I (SWITCH), Z (STOP)
           P (PURGE),  R (REPLICATE), H (HELP),   Q (QUIESCE LOGGING)       <<<
 LC  PARM     TITLE                     USER ID  TARGET TYPE AREA  STAT  LOG
     PAYROLL  PAYROLL TRANSACTIONS      CIR11    IMSA    DET IWKLD ACTV  ACT
     APSB01   TRACE OF APSP01           PWW1     IMSA    DET IWKLD ACTV  QIS
     JSC      WORKLOAD TRACE            JSC4     IMSA    SUM IWKLD ACTV
  ****************************** END OF REQUESTS ****************************
```

Figure 20. Current Traces Application

Trace requests are shown in the list as they are processed. This application allows service selection by line command and shows all the active requests, parameters specified for each request, the user logon identification, the target IMS of the request, the service security classification, the area of IMS being monitored, the service status, and the logging status for this request.

Selecting an active trace request invokes the LTRAC service for that trace. Other line commands can be used to view current trace request options, access data entry panels to replicate or change current options to make another unique trace request, or confirm a purge of the selected request.

The fields for the Current Traces application are:

Field Name	Description
INTVL	Screen refresh field. You can specify a refresh interval for this application. Entering GO in the COMMAND field or pressing PF6/18 starts the refresh cycle.
LC	A line command input field. One-character line commands can be entered in this field to view, modify, or replicate the options for a selected request; purge a request; or display HELP information about the service (see “Line Commands” on page 244).
PARM	This field shows the parameters defined for the active requests.
TITLE	The service title.

USER ID	The logon identification of the user who made the request.
TARGET	The IMS subsystem defined as the target of the requested service either by default or user-specified.
SEC	The security code for user access to the service.
AREA	The IMS resource area being analyzed. This field could contain:

Field Data	Description
DB	Database activity and buffer pool utilization
IMVS	IMS and OS/390 interactions
INTNL	IMS internal functions
IWDB2	IMS workload DB2 activity
IWFP	IMS workload Fast Path activity
IWGBL	IMS workload global IMS region calls
IWKLD	IMS workload
IWTRN	IMS workload transactions
LOCK	IRLM functions
MFS	Terminal I/O
QUEUE	IMS queuing
REGN	Application program activity in the dependent regions
SCHED	Scheduling of application programs in the dependent region

STAT	The service request status, which could be:

Field Data	Description
ACTV	The request is active.
COMP	The request executed and completed normally.
HELD	The request is being held and is pending release.
INIT	The request is being invoked for the first time (a start time was specified, but it has not been reached).
INV	The request terminated because of an invalid parameter or measurement. The BBI-SS PAS Journal log contains a descriptive message of the error.
LOCK	A LOCK command was issued for the service or the service abended.
QIS	The service is quiesced, because the target IMS is not active.

	RST	The target IMS restarted. The request is waiting until the current interval expires before performing restart processing as specified by the RST keyword in the original request.
LOG	The trace logging status for this request. Possible values:	
	blank	Logging was not requested for this trace.
	ACT	Logging is active.
	QIS	Logging is quiesced.
	SUSP	Logging is suspended because the trace is quiesced. Logging resumes when the trace restarts.

Application Transfer Commands

The following related application transfer commands can be entered in the COMMAND field of the Current Traces list application showing all active trace requests:

HT (HISTORY TRACES)
ST (START TRACES)

HT (HISTORY TRACES)

HT displays the History Traces application (see Chapter 31, “Displaying History Traces (Primary Menu Option 4)” on page 247), which is used to manage trace log data sets where traces are recorded. Use this command to access and control the current and historical trace logs.

ST (START TRACES)

ST displays the MTRAC data entry panel to start a trace request. Specify the keyword parameters, selection criteria, and exception filters in successive panels to activate data collection for a summary or detail trace as described in Chapter 27, “Requesting Workload Trace Data Collection (MTRAC)” on page 203.

TYPE Primary Command

Entering TYPE in the COMMAND line lists the traces of the specified type only. The possible types that can be specified are shown in the TYPE column of the History Traces application. For example, to list only detail traces), enter:

```
COMMAND ===> TYPE DET
```

Enter TYPE to return to the list of all traces.

Line Commands

Entering one of the following one-character line commands in the LC field for a request executes the line command function. Multiple selections can be entered at one time by selecting a series of services and pressing the ENTER key. Each display in a series is processed by pressing the END key. Each data entry timer request panel in a series that is to be modified or purged is submitted by pressing the ENTER key and then the END key to process the next request.

Line Command	Description
S	SELECT. Displays a list of scrollable trace entries. Selecting an active workload trace (MTRAC) request displays the collected IMS trace entries (LTRAC) for the selected trace request.
W	SHOW. Shows a display panel of the trace request options defined for the selected request for viewing only (see “Show Trace Options Request (W Line Command)” on page 246).
M	MODIFY. Shows a data entry panel of the trace options defined for the selected request so the options can be changed (see “Modify Trace Options Request (M Line Command)” on page 245).
I	SWITCH. Switches logging of traces from the current VSAM data set to the next available one (see “Switch Log Data Set Request (I Line Command)” on page 246).
Z	STOP. Stops the request and retains collected data. The STOP time equals the current time. **Note:** You must be authorized to issue the Z (STOP) line command for traces other than your own. This is done with the PMACC parameter in a user authorization member of the BBPARM data set as described in *Implementing Security for MAINVIEW Products*.
P	PURGE. Displays a PURGE panel to verify a purge of the selected request (see “Purge a Trace Request (P Line Command)” on page 245).
R	REPLICATE. Shows data entry panels for the trace options defined for the selected request so that the options can be repeated or changed to make a new request for that service (see “Replicate a Trace Request (R Line Command)” on page 246). The request must be unique (defined by the service select code plus a parameter).
H	HELP. Displays HELP information about the service for the selected request.
Q	QUIESCE LOGGING. Quiesces logging of traces to an external VSAM data set (see “Quiesce Trace Logging Request (Q Line Command)” on page 245).

List of Trace Transactions Display (S Line Command)

The S (SELECT) line command for a trace displays the LTRAC service. LTRAC shows a scrollable list of transaction instances for the selected trace. From here, all other trace displays can be accessed with EXPAND as described in Chapter 34, “LTRAC - List of Trace Entries” on page 265.

Modify Trace Options Request (M Line Command)

Selecting a trace request with the M line command displays a data entry panel(s) with the options that were defined to activate the selected trace (see Chapter 27, “Requesting Workload Trace Data Collection (MTRAC)” on page 203). Only those options with an ===> can be changed.

Note: If trace logging to external VSAM data sets was specified with the LOGTRAC option (Y for yes to start trace logging; N for NO to stop it), it cannot be changed with the modify command while a trace is active. For more information about trace logging and the options you can change after logging is active, see “Modify a Trace Log Request (M Line Command)” on page 236.

Purge a Trace Request (P Line Command)

Selecting an active trace request with the P line command displays a purge confirmation panel like the one shown on page 74. Pressing the ENTER key confirms a purge of the selected service request. A short message in the upper right corner of the display shows the result of the request. If an ERROR IN REQUEST message is displayed, a short explanatory message is displayed on the third line. Pressing the END key (PF3/15) redisplays the Current Traces list application.

To stop a request and retain a trace, use the modify command and enter a stop time. This allows information to remain after collection stops. If a request is purged, all data is lost.

Note: You must be authorized to issue the P (PURGE) line command for traces other than your own. This is done with the PMACC parameter (refer to *Implementing Security for MAINVIEW Products).*

Quiesce Trace Logging Request (Q Line Command)

If the Current Traces application shows ACT for the LOG field of a request, the trace is being logged to external VSAM data sets (see Chapter 28, “Logging a Trace” on page 221). Selecting that trace request with a Q line command immediately quiesces logging of the trace.

Once quiesced, logging cannot be restarted for this request. The trace continues to run, but no data is written to the log. To restart logging, the request must be purged with the P line command. Another request for this trace with the logging options specified must be activated.

Note: You must be authorized to issue the Q (QUIESCE LOGGING) line command for traces other than your own. This is done with the PMACC parameter (refer to *Implementing Security for MAINVIEW Products).*

Replicate a Trace Request (R Line Command)

The R line command displays the MTRAC Start IMS Trace Request data entry panel and each of the additional MTRAC options panels, described in Chapter 27, “Requesting Workload Trace Data Collection (MTRAC)” on page 203, for the selected trace request. Each of the panels shows the options specified when the original trace request was activated. All values prefixed with an ===> can be modified.

Note: If the request specified a STOP time, it appears as nnnn, where nnnn is the number of intervals remaining until the trace stops.

The R line command is used to replicate a request to define another one. Each request must be unique with an ID in the PARM field that identifies the request. The display is preset with the service select code.

Pressing the ENTER key submits the request. A short message in the upper right corner of the display shows the result of the request. If an ERROR IN REQUEST message is displayed, a short explanatory message is displayed on the third line. Pressing the END key (PF3/15) redisplays the Current Traces list application.

Show Trace Options Request (W Line Command)

Selecting a trace request with the W line command shows the MTRAC Start IMS Trace Request panel and each of the additional MTRAC options panels, if requested, used to submit the selected trace (see Chapter 27, “Requesting Workload Trace Data Collection (MTRAC)” on page 203). If trace log options were specified with the trace request, additional fields are shown with the previously defined trace log options. These fields are obtained from the current trace log data set (see “Show a Trace Log Request (W Line Command)” on page 235).

Each option is suffixed by a colon (:), which means the option value cannot be changed. The ===> for the “additional trace options” fields allows you to enter a Y if you want to review any additional trace options you have defined previously. Press the END key to return to the Current Traces list application.

Switch Log Data Set Request (I Line Command)

If a trace is being logged to external VSAM data sets (see Chapter 28, “Logging a Trace” on page 221), the Current Traces application shows ACT for the LOG field of the trace request. You can request that trace logging be switched immediately to the next available data set by selecting that trace request with the I line command. An archive request for the trace log data set just deactivated is started if the selected MTRAC request specified the archive option.

If the switch command is issued and only one log data set is defined, logging either reuses the same data set or quiesces. Logging resumes with the same data set if the Overwrite Log option for the MTRAC request was Y (yes). Logging quiesces if the Overwrite Log option for the MTRAC request was N (no).

Note: You must be authorized to issue the I (SWITCH) line command for traces other than your own. This is done with the PMACC parameter (refer to *Implementing Security for MAINVIEW Products).*

Chapter 31. Displaying History Traces (Primary Menu Option 4)

This application is provided by the Workload Analyzer. It is used to display historical traces recorded to trace log data sets and to manage these data sets. The recording of traces to trace log data sets can be defined with an MTRAC service request to activate a trace. To use this History Traces application, you should understand how to request trace logging and how it works as described in Chapter 27, “Requesting Workload Trace Data Collection (MTRAC)” on page 203 and Chapter 28, “Logging a Trace” on page 221.

History Traces can be accessed by:

- Selecting Primary Menu Option 4 from the Primary Option Menu
- Entering an HT command to transfer to this application from another one.

```
BMC SOFTWARE  ---------------- HISTORY TRACES  --------------  PERFORMANCE MGMT
 COMMAND ===>         TGT ===> IMSA
                                            TIME--  09:36:34  SCROLL ===> CSR
 COMMANDS: SORT, LOCATE, NEW, STOP, START, TYPE
 LC CMDS:  S (SELECT), W (SHOW), P (PRINT)   D (DELETE),  E (RESET)
           V (VERIFY), N (NEW),  A (ARCHIVE), F (FREE)

  DIRECTORY:    CIR4.LL1X.TRACEDIR
  ENTRIES USED: 1,209     FREE:   201
                                                              SCROLL  RIGHT >>>
 LC  DATE------TIME TRACEID  TITLE                   USERID   TARGET   STAT ACT
     00/12/15 22:00 TEST001  TEST DETAIL TRACE       CIR4     IMSA     INV
     00/12/15 21:00 DETAIL   WORKLOAD DETAIL TRACE   CIR2     IMSA     USD  RD
     00/12/15 00:00 BLNK     TRACE OF SPECIAL TRANS  BABUSERS IMSA     UPD  WRT
     **************************** END OF DATA **********************************
```

Figure 21. History Traces Application (Before Scrolling Right)

```
BMC SOFTWARE  ----------------- HISTORY TRACES --------------  PERFORMANCE MGMT
 COMMAND ===>         TGT ===> IMSA
                                            TIME--  09:36:34  SCROLL ===> CSR

 COMMANDS: SORT, LOCATE, NEW, STOP, START, TYPE
 LC CMDS:  S (SELECT), W (SHOW), P (PRINT),   D (DELETE),  E (RESET)
           V (VERIFY), N (NEW),  A (ARCHIVE), F (FREE)

  DIRECTORY:    CIR4.LL1X.TRACEDIR
  ENTRIES USED: 1,209     FREE:   201
                                                                SCROLL LEFT <<<
 LC  DATE------TIME  TYPE  DSN                                         VOLSER
     00/12/15 22:00  DET   CIR4.IMSA.LEOTST01.NOV28.T2200.V01          MIGRAT
     00/12/15 21:00  DET   CIR21234.IMSA.DETAIL1X.TEST1234.V02         BAB001
     00/12/15 00:00  SUM   CIR2.IMSPROD.DB2E.TRANS.NOV01.T0000.V01     BAB303
     **************************** END OF DATA **********************************
```

Figure 22. History Traces Application (After Scrolling Right)

The History Traces application displays the trace directory which tracks all trace logs. You can use this application with simple line and primary commands to:

- Maintain the trace directory
- See an online display of historical trace data records whether the trace was recorded the same day, last night, a year ago, or on another system
- Review trace options
- Do online administration of the trace log data sets (purge, free, verify, reset, archive)
- Move trace log data sets from one system to another

This application displays the trace directory showing one row for each log data set tracked in the directory. The initial display is in descending order of date and time but can be sorted by any column. For example, if you want to see all the data sets for one trace together, you can sort the list by TRACEID.

This display contains the following:

Field Name	Description
DIRECTORY	Data set name of the directory.
ENTRIES USED	Number of entries in the directory that are used.
FREE	Number of entries in the directory that are not used. This field is highlighted if fewer than 10 entries remain. **Note:** You can use the batch maintenance jobs described in Chapter 28, "Logging a Trace" on page 221 to clean up or expand the directory.

The fields for the list of trace log data sets are described below in alphabetical order.

ACT	Status of log data set activity: blank — Data set is not currently in use. ALL — Data set is allocated to a currently active trace. R/W — One or more users are reading the data set and an active trace is also writing to it. REA — One or more users are reading the data set, but nothing is being written to it. WRI — An active trace is writing to the data set.
DATE----TIME	Date and time that logging was activated on this data set.
DSN	Name of the trace log data set.

STAT	Status of the trace log data set contents:		
	EMP	Data set contains no trace records.	
	INC	Data set did not close or archive successfully after an update; for example, in a system failure.	
	INV	Data set is an invalid trace log data set. The exact reason is contained in a message written to the BBI-SS PAS journal.	
	NOC	Data set is not cataloged.	
	RES	Data set contains valid trace data and is flagged for reuse.	
	UPD	Data set is being updated.	
	USD	Data set contains valid trace data and closed successfully.	
TARGET	Target that was specified when the trace was requested.		
TITLE	Title specified when the trace was requested.		
TRACEID	Trace identifier specified when the trace was requested.		
TYPE	Type of trace, either SUM (summary) or DET (detail), specified when the trace was requested.		
USERID	ID of the user who requested the trace.		
VOLSER	Volume serial name of the first volume used for this data set.		

SORT Primary Command

When the list of trace log data sets is displayed initially, the list is sorted in descending order by date and time. Any of the columns in the display can be sorted with the SORT command. The sort, except for date and time, is in ascending order. SORT is entered in the COMMAND field of the display with the first two characters of the column name as:

```
SORT cc
```

where cc can be any of the following two characters, which are described below in alphabetical order.

AC Sorts the list by activity status of the log data set.

DA Sorts the list by the date and time logging was activated on this data set.

Note: DATE-TIME is treated as one field so you cannot sort by time without also sorting by date.

DS Sorts the list by data set name of the trace log.

ST Sorts the list by status of the trace log data set contents.

TG Sorts the list by name of the target system.

TI Sorts the list by trace title.

TR Sorts the list by trace identifier.

TY Sorts the list by type of trace (detail or summary).

US Sorts the list by user ID.

VO Sorts the list by volume serial name.

LOCATE Primary Command

Entering LOCATE in the COMMAND line moves an entry to the top of the log data set list by the currently sorted column. It is similar to the ISPF LOCATE member command.

NEW Primary Command

Entering NEW in the COMMAND line presents a data entry panel, shown in Figure 23 on page 251, to manually add trace log data sets to the directory; for example, to move them from one system to another. Enter the data set name only. The application automatically verifies each name entered.

Figure 23 shows a sample data entry panel that specifies a data set to be added to the directory. Pressing ENTER processes the request and shows the results in the RESPONSE field, which can be:

CANNOT ALLOCATE
: Trace log data set cannot be allocated.

CANNOT READ
: A non-VSAM linear data set was specified that is not a valid data format for a trace log data set.

DOES NOT EXIST
: Data set name is not cataloged.

NOT TLDS
: Request to add a trace log data set (TLDS) to the directory could not be processed, because the work area could not be obtained or the data set does not contain trace data.

TLDE ALREADY EXISTS
: Trace log data set already exists in the directory.

TLDS ADDED
: Trace log data set added to the directory.

```
BMC SOFTWARE  ---------- CREATE TRACE DSN ENTRIES -----------  PERFORMANCE MGMT
COMMAND ===>        TGT ===> IMSA
                                                            TIME --  09:36:34

  DIRECTORY: CIR4.LL1X.TRACEDIR

  SPECIFY TRACE DATASET NAME(S), ONE NAME PER ROW.
  THE HEADER RECORD IN EACH DATA SET WILL BE READ TO CREATE THE DIRECTORY ENTRY.

         LOG DATA SET NAME                              RESPONSE
        --------------------------------------------   -----------------------
    1.  CIR4.LL1X.IMSA.TEST001.MAY21.T2200.V02         TLDS ADDED
    2.
    3.
    4.
    5.
    6.
    7.
    8.
    9.
    10.

 Press ENTER to add datasets
```

Figure 23. New Primary Command Data Entry Panel

STOP Primary Command

Entering STOP in the COMMAND line closes and deallocates the trace directory. Active trace logging continues. All commands except START are disabled and a new trace logging request cannot be initiated. A trace directory is required for trace logging to occur.

Note: This is done with the PMACC parameter (refer to *Implementing Security for MAINVIEW Products).*

START Primary Command

Entering START in the COMMAND line reallocates the directory and opens it with a disposition of old. If this is successful, all commands are re-enabled and new trace logging requests can be initiated. If START is not successful, all other application commands remain disabled and trace logging cannot be initiated.

Note: This is done with the PMACC parameter (refer to *Implementing Security for MAINVIEW Products).*

TYPE Primary Command

Entering TYPE in the COMMAND line lists the traces of the specified type only. The possible types that can be specified are shown in the TYPE column of the History Traces application. For example, to list only detail traces, enter:

```
COMMAND ===> TYPE DET
```

Enter TYPE to return to the list of all traces.

Line Commands

Entering one of the following one-character line commands in the LC field for a log data set executes the line command function. Multiple selections can be entered at one time by selecting a series of services and pressing the ENTER key. Each display in a series is processed by pressing the END key. Each data entry request panel in a series is submitted by pressing the ENTER key and then the END key to process the next request.

Note: If your user ID does not match the user ID of the trace, you must be authorized to use the following line commands:

- Delete
- Free
- Reset
- Archive

Note: This is done with the PMACC parameter (refer to *Implementing Security for MAINVIEW Products).*

Line Command	Description
S	SELECT. Loads the selected log data set and displays the LTRAC screen. From there, you can access all other trace displays, as described in Chapter 34, “LTRAC - List of Trace Entries” on page 265.
W	SHOW. Displays the selected log data set information. The directory entry is updated if necessary.
P	PRINT. Generates the print JCL you can submit to print trace data from the log data set. **Note:** This command is available only when the terminal session is operating under ISPF.
D	DELETE. Deletes the entry from the trace directory. If the data set is cataloged, a confirmation screen similar to the sample in Figure 24 is displayed. If the entry is for a data set that is not cataloged, no confirmation is required.

```
BMC SOFTWARE  -------- TRACE DIRECTORY PURGE CONFIRMATION ---- PERFORMANCE MGMT
COMMAND ===>                                                    TGT ===> IMSA
                                                             TIME --  09:36:34

 LOG DATA SET NAME:   CIR4.LL1X.TEST1.V02
 TARGET:              IMSA
 USERID:              CIR4
 TRACEID:             BIGELAP
 VOLUME:              ARCHIV
 START DATE-TIME:     00/12/15 22:00
 STATUS:              INV

 Do you wish to also delete the log data set? ===> Y  (Y/N)

 Press ENTER to confirm; END to cancel
```

Figure 24. Trace Directory DELETE Entry Confirmation Panel

E	RESET. Marks the log data set for reuse.
V	VERIFY. Matches the directory against the log data set and updates the directory with information from the log data set. If the data set is not cataloged, the status is changed to NOCAT. If there is any error in reading this information, the status is changed to INV (invalid).
N	NEW. Acts like the NEW primary command, except the data entry screen is primed with the name of the selected log data set.
A	ARCHIVE. Initiates the archive started task (STC) for the selected log data set if the archive option was specified with the MTRAC trace request. If archiving was not specified for this log data set, the request is ignored.
F	FREE. Closes and deallocates the selected log data set if a trace is not being written to it. This command also can be used when another user is viewing the same trace log.

Chapter 32. Printing History Traces (Primary Menu Option 4)

This chapter describes how to print history trace data using either the offline or the online printing facility.

Primary Input - Trace Log Data Sets (TLDSs)

The primary input source is data that MVIMS collects and stores in trace log data sets. The normal input is the VSAM trace log data set (a TLDS not archived to tape) created by a trace request with history logging active at the time of the trace. All trace log data sets (TLDSs) are tracked in a trace directory and can be accessed through the History Traces application.

You can specify one or more TLDSs in a series of DD statements with the ddnames in the form TRACINxx. Only TLDS data sets can be specified in this manner.

If you don't know the exact names of the trace data sets, you can use the TRACEID keyword to specify which trace is to be printed. In this case, you must also use the TRACEDIR DD statement to specify the directory data set. It is recommended but not required that you further qualify which trace data sets are to be selected by using the TARGET, DATE, and TIME keywords.

Alternative Input - Archived TLDSs

An archived TLDS is produced by the archive utility IMFTARC. To print directly from this sequential data set without reloading it to a VSAM linear data set, specify the DSN on the ARCIN DD statement.

Output Formats

The output format is the same as that of the online displays; it is 90 characters wide with the first position reserved for the attribute character.

Printing from an Online Application

MVIMS allows you to print trace log data sets online from the History Traces application if you are executing your TS from ISPF. You can do this by using the P (PRINT) line command.

Before printing from the online application, you must first copy the skeleton JCL located in the WATBPRNT member of the BBPROF data set that is distributed as part of the MVIMS-distributed library.

Copy WATBPRNT to an individual user data set (BBPROF) or to a site data set (SBBPROF). The BBPROF or SBBPROF data set must be defined in the CLIST (MAINVIEW CLIST) used to start the terminal session. For more information about BBPROF, see the *MAINVIEW Common Customization Guide*.

To print from the online application:

1. Select Option 4 from the MVIMS Main Menu.
2. In the History Traces panel, enter the P line command next to the data set you wish to print. (For a full description of how to use line commands in the History Traces application, see “Line Commands” on page 252 in Chapter 31, “Displaying History Traces (Primary Menu Option 4)”.)
3. Enter the required information in the next panel that appears. This panel gives you options that allow you to tailor the print job output to your needs.
4. Press the End key.

The printed trace data has the same format and content as the online displays.

Printing Using a Batch Utility Job

MVIMS provides you with the option of printing history trace data offline. You can do this by using a batch utility job, WATBTRAC. Before submitting the WATBTRAC batch print job, you must tailor the JCL described in the next section.

JCL

In order to print trace log data sets offline, you must modify and then submit the JCL described in this section. The skeleton JCL to print trace data sets is shown in Figure 25. This JCL is a single-step procedure located in member WATBTRAC of the BBSAMP data set.

```
//         JOB (ACCT),'NAME'
//WATBPRNT PROC TLDS=NULLFILE,           INPUT TRACE DATA SET
//             TDIR=NULLFILE,            INPUT TRACE DIRECTORY
//             ARC=NULLFILE,             INPUT ARCHIVED TRACE DATA SET
//             PFX='HILVL.RUN.LIB'       DSN PREFIX OF BBLINK
//PRINT   EXEC PGM=WATBPRNT,REGION=4M,PARM='GMWK=128K'
//*                                      INCREASE GMWK FOR LARGE TRACES
//STEPLIB   DD DISP=SHR,DSN=&PFX..BBLINK
//SYSPRINT  DD SYSOUT=*                  INPUT LIST AND DIAGNOSTICS
//SYSUDUMP  DD SYSOUT=*                  ABEND DUMPS
//STD1      DD SYSOUT=*                  DEFAULT REPORT OUTPUT
//*
//*    ONLY ONE OF THE FOLLOWING INPUT DD'S MAY BE SPECIFIED
//*
//TRACIN01  DD DISP=SHR,DSN=&TLDS        TLDS INPUT
//*TRACEDIR  DD DISP=SHR,DSN=&TDIR       TLDS INPUT THRU TRACE DIR
//*ARCIN     DD DISP=SHR,DSN=&ARC        ARCHIVED TLDS INPUT
//        PEND
//*
//*       SPECIFY INPUT FILE HERE AS TLDS=
//*
//PRINT   EXEC WATBPRNT,TLDS='SYS5.IMSP.THRDHIST.JUL01.T0001.V01'
//REPORT1   DD SYSOUT=*                  USER-DEFINED OUTPUT DD
//SYSIN     DD *                         CONTROL CARD INPUT
    REPORT  REPORTID=WATBPRNT,
            DDNAME=REPORT1,
            LTRAC=YES,
            STRAC=YES,
            DTRAC=YES
```

Figure 25. JCL to Print a Trace (WATBTRAC)

The individual JCL control statements, shown in Figure 25, specify how to read the input data set for each of the requested traces and how to format the output. The statements are described below.

PRINT

Specifies the name of the program as:

`PGM=WATBPRNT`

and the region required to run the program.

This program supports a `GMWK` option in the PARM field. For example:

`PGM=WATBPRNT,PARM='GMWK=128K'`

This option is used to increase the size of the summary work area needed to process trace data sets that contain a very large number of transactions. This option increases the summary work area for all services (LTRAC, STRAC, and so on).

The syntax rules for GMWK are similar to those for the OS/390 JCL REGION parameter. The size specified should be greater than 128K. Storage acquired with this option is below the 16M line and is limited by the largest private area available in the OS/390 system. When the GMWK option is used, the job's REGION size may need to be adjusted accordingly.

A general guideline for the size of GMWK is:

```
GMWK = 1700 + (72 * number of unique transactions
summarized)
```

STEPLIB DD — Defines the program library containing the WATBPRNT load module.

SYSPRINT DD — Defines the data set for printing all input statements and program messages.

Note: The SYSPRINT DD statement must exist. If it is not found, a WTO is issued and the run is terminated.

SYSUDUMP DD — Defines the dump data set for problem determination.

STD1 — Defines the ddname for the default report.

TRACINxx DD — Defines one or more trace log data sets (TLDSs) as input for trace data formatting. If TRACINNxx is defined, TRACEDIR and ARCIN are ignored. This DD must refer to a VSAM trace data set (a TLDS not archived to tape), written by the online trace.

Multiple TLDS data sets can be processed by specifying each in a DD statement with a ddname of TRACINxx, where xx is any valid set of characters. Each of the TLDSs is processed in the order it exists in the job stream.

TRACEDIR DD — Designates a trace directory data set.

If no TRACINxx DD statement is found, the TRACEDIR DD statement is processed.

If you specify TRACEDIR DD, you must use the TRACEID keyword. All trace data sets defined with this trace ID are eligible for printing.

ARCIN DD — Defines an archived TLDS. TRACIN and TRACEDIR must be dummied or you must specify NULLFILE.

REPTDD1 — Defines the ddname for the trace print output.

SYSIN DD — Defines the input data set containing the control card specifying the amount of detail to print.

Request Keywords

This section identifies keywords that can be used to request a trace print.

Table 13. Identification Keywords

<table>
<tr><th>Keyword/Operand</th><th>Description</th></tr>
<tr><td>AB|ABEND</td><td>Selects only those transactions with a nonzero system or user completion code.</td></tr>
<tr><td>DATE=[ddmmmyyyy|
ddmmmyyyy-ddmmmyyyy]</td><td>The date value is used to select a date range or starting date from the input. The input can be a single date in ddmmmyyyy format, which defines a starting day, or a range in the form ddmmmyyyy-ddmmmyyyy to define starting and ending dates to process.

The asterisk is used to specify current date. To select all entries from yesterday, specify:

<code>DATE=*_1</code>

Note: The underscore is used instead of a dash, since the dash is already used to define a range of dates. Using *-1 would result in interpreting the date as a range from current date to day 1 (which is invalid as a date).

To process only one day, specify:

<code>15DEC2000-15DEC2000</code>

The default is to select all records.</td></tr>
<tr><td>DB2C|DB2CNT<nnnnnn|>nnnnnn|
=nnnnnn</td><td>Where nnnnnn represents number of DB2 database calls. It selects a subset of the trace entries in the online trace storage buffer or trace log data set by DB2 database call counts.</td></tr>
<tr><td>DB2T|DB2TIME<nnnnnn|>nnnnnn|
=nnnnnn</td><td>Where nnnnnn represents CPU time in milliseconds. It selects a subset of the trace entries in the online trace storage buffer or trace log data set by DB2 elapsed time.</td></tr>
<tr><td>D|DAY=nn|nn-nn</td><td>Where nn represents a day or range of days as a relative number from the date of the first trace entry (first date is 1, up to a maximum of 99) in the trace buffer or trace log data set. It selects a subset of the entries by a day or range of days.

If there is only one day's data, the END TIME timestamp is hh:mm:ss.th. If the trace buffer or trace log data set contains more than one day's data, END TIME changes to dd-hh:mm:ss, where dd is the day number relative to the date of the first record in the trace buffer or trace log data set.

If DAY is not specified, but TIME is, the default is the day of the most current record in the buffer or trace log data set.</td></tr>
<tr><td>DLIC|DLICNT<nnnnnn|>nnnnnn|
=nnnnnn</td><td>Where nnnnnn represents number of DL/I database calls. It selects a subset of the trace entries in the online trace storage buffer or trace log data set by DL/I database call counts.</td></tr>
</table>

Table 13. Identification Keywords (continued)

Keyword/Operand	Description
DLIT|DLITIME<nnnnnn|>nnnnnn|=nnnnnn	Where nnnnnn represents DL/I CPU time in milliseconds. It selects a subset of the trace entries in the online trace storage buffer or trace log data set by DL/I elapsed time.
DTRAC=[YES|NO]	The DTRAC keyword requests printing a one-line entry for each detail trace event. For ease of associating the detail data with a transaction, it is recommended that either LTRAC or STRAC also be requested, but this is not required. The possible values are YES and NO. The default is DTRAC=NO.
EL|ELAPSED<nnn|>nnn|=nnn	Where nnn represents elapsed time in milliseconds. It selects a subset of the trace entries in the online trace storage buffer or trace log data set by elapsed time.
IOC|IOCNT<nnnnnn|>nnnnnn|=nnnnnn	Where nnnnnn represents number of total I/Os (READs and WRITEs). It selects a subset of the trace entries in the online trace storage buffer or trace log data set by total I/O count.
LTRAC=[YES|NO]	The LTRAC keyword is used to request the LTRAC one-line entry for each transaction. If used in conjunction with STRAC or DTRAC, each transaction starts with the LTRAC line followed by the other displays as requested. The possible values are YES and NO. The default is LTRAC=YES.
MSW	Selects only those transactions that are a result of a message switch.
PR|PROGRAM=xxxxxxxx	Where xxxxxxxx represents a program name. It selects a subset of the trace entries in the online trace storage buffer or trace log data set by program name. Multiple trace entries can be selected by using an * or + qualifier; for example, `PR=SM++0+00`.
REPORTID={name}	Each trace print must have a unique identification provided by REPORTID. This is the only required keyword. REPORTID is also used as the ddname of the output data set. The value can be 1–8 characters, with any special characters allowed.
RE|RESPONSE<nnnnnn|>nnnnnn|=nnnnnn	Where nnnnnn represents response time in milliseconds. It selects a subset of the trace entries in the online trace storage buffer or trace log data set by response time.
STRAC=[YES|NO]	The STRAC keyword is used to request the STRAC transaction summary display. STRAC=NO suppresses the STRAC displays. STRAC=YES requests only the first section of the STRAC display. These requests can be combined with any of the DTRAC requests. The default is STRAC=NO.

Table 13. Identification Keywords (continued)

Keyword/Operand	Description
T|TIME=`hhmm|hhmm-hhmm`	Where `hhmm` represents a start time or time period. It selects a subset of the trace entries in the online trace buffer or trace log data set by a start time or time period. When the start time is higher than the end time, for example, `TIME=2000-0300`, 24 hours are added.
TITLE1=[c.....c]	The title value is printed centered on the first line of each page of the report. The value can be up to 60 characters. If this keyword is not included, the title is left blank on the report.
TITLE2=[c.....c]	The secondary title is printed centered on the second line of each page of the report. The value can be up to 60 characters. If this keyword is not included, the secondary title is taken from the title specified when the trace was created.
TRACEID=[id|BLANK]	Defines the ID of the trace that is to be matched in a search through the trace directory. All TLDS data sets that meet this requirement, plus any TARGET, TIME, and DATE requirements, are dynamically allocated and processed. To define a trace ID of blank, specify: `TRACEID=BLANK` **Note:** The DD statement TRACEDIR must be present before this keyword is processed. If the DD statement is not present, this keyword is ignored.

Table 13. Identification Keywords (continued)

Keyword/Operand	Description
TY|TYPE=xxx	Where xxx represents a region type, which can be: BMP — Batch message processing or Java batch message processing (JBP) region. CON — Conversational MPP or JMP region. DBT — DBCTL thread (CICS and ODBA). FPU — Fast Path utility region. MDP — Message-driven Fast Path region. MPP — Message processing or Java message processing (JMP) region. NDP — Non-message-driven Fast Path region. TPI — Message processing region currently executing an explicit CPI-C program. It selects a subset of the trace entries in the online trace storage buffer or trace log data set by specified region type. Multiple trace entries can be selected by using an * or + qualifier; for example, `TY=M*`.
US|USER=xxxxxxxx	Where xxxxxxxx represents a user or LTERM ID. It selects a subset of the trace entries in the online trace storage buffer or trace log data set by user or LTERM. Multiple trace entries can be selected by using an * or + qualifier; for example, `US=CICS*`.

Chapter 33. Workload Analyzer Trace Services (Quick Reference)

The following table shows the service select code and parameters for the Workload Analyzer trace services. It also lists the page where you can find a complete description of each service.

Table 14. Workload Trace Service Select Codes

<table>
<tr><th>Service Select Code</th><th>Parameter</th><th>See</th></tr>
<tr><td>MTRAC</td><td>[traceid]
[keywords]</td><td>Chapter 27, “Requesting Workload Trace Data Collection (MTRAC)” on page 203</td></tr>
<tr><td>LTRAC</td><td>[,traceid
,AB
,D
,DB2C<nnn|>nnnn|=nnn
,DB2T<nnn|>nnnn|=nnn
,DLIC<nnn|>nnnn|=nnn
,DLIT<nnn|>nnnn|=nnn
,EL<nnn|>nnnn|=nnn
,IOC<nnn|>nnnn|=nnn
,MSW
,PR=xxxxxxxx
,RE<nnn|>nnnn|=nnn
,T=hhmm|hhmm-hhmm
,TR=xxxxxxxx
,TY=xxx
,US=xxxxxxxx]</td><td>Chapter 34, “LTRAC - List of Trace Entries” on page 265</td></tr>
<tr><td>STRAC</td><td>[traceid|,SEQ=n]</td><td>Chapter 35, “STRAC - Summary Trace Data Display” on page 273</td></tr>
<tr><td>DTRAC</td><td>[traceid|,LEVEL=1|2|3|4
,SEQ=n]</td><td>Chapter 36, “DTRAC - Detail Trace Data Display” on page 287</td></tr>
</table>

Chapter 34. LTRAC - List of Trace Entries

```
BMC SOFTWARE  ------------ IMS TRACE ENTRIES-------------------PERFORMANCE MGMT
SERV ==> LTRAC        INPUT    11:44:25  INTVL=> 3  LOG=> N     TGT==> IMSxxx
PARM ==>                                         ROW 1 OF 506    SCROLL=> CSR
EXPAND:  MON(WKLD), HISTORY, LINESEL(STRAC)  ENTRIES IN DATASET 45 - 550
                                                           HIST TGT==> IMSxxx
 31DEC94
 END TIME    PROGRAM  TRANCODE  TYP RESPONSE  ELAPSED      DLI      DB2 USR/LTRM
-----------  -------- --------  --- -------- -------- -------- -------- --------
09:20:23.12  PIWAIT   TDRIVER   BMP      50s      50s       20        0 CIR11
09:20:22.32  PHDAMINQ THDAMINQ  MPP  3,330ms  3,230ms      200       90 PWW1
09:20:12.04  PHDAMINQ THDAMINQ  MPP  8,245ms* 3,450ms      100       10 PWW1
09:09:13.90  PHIDMINQ THIDMINQ  MPP    <10ms    <10ms        7       10 PWW1
```

Description: This service provides a scrollable list of all the active trace entries in one of the following:

- The online trace storage buffer for an active MTRAC request selected from the Current Traces application (Primary Menu Option 3)
- A trace log data set selected from the History Traces application (Primary Option Menu 4) Each trace entry represents a transaction instance.

LTRAC can be used to see all occurrences of a transaction with the timestamp of when each completed. You can review this list for application performance and select any one of these occurrences to view further detail with the STRAC display.

Color: If you have a color monitor:

Turquoise	Indicates normal values.
Red	Highlights any transactions that have abended.
White	Indicates column headings.
Yellow	Indicates error messages.

Select Code: LTRAC

Parameter: All of the LTRAC parameters act as filters that restrict the information shown according to the criteria specified by the parameter(s). The LTRAC parameters can be used as follows:

- The filtering parameters limit the data shown by the specified parameter.

  For example, if you enter:

  `TY=DBT`

  only the rows for DBCTL CICS and ODBA threads are displayed.

- Multiple parameters must be separated by commas.
- A blank indicates the end of a parameter string.
- Multiple resources with similar names can be requested by using an * character as a generic qualifier and a + character as a positional qualifier. The positional qualifier is repeated for every character to be replaced. The generic qualifier replaces all characters at the end of the parameter.

  For example, a parameter of `NAME=MP*` shows all region names that start with MP and have any character in the rest of the name.
- If multiple filtering parameters are entered, the data displayed must meet **all** the restrictions.
- If one parameter invalidates another, an error message is issued without further processing.

The following parameter descriptions are arranged alphabetically except for `traceid` which is a positional parameter and must always be in the first position. The descriptions begin with parameters that start with a special character. Parameters containing a numeric character are first in their alphabetical group.

`traceid`

Where `traceid` is a unique ID used with the MTRAC request. A blank can be an identifier. This parameter is positional and must be in the first position. This field is primed automatically with the MTRAC identifier if the request is selected from the Current Traces or Active Timer requests application.

AB|ABEND

Selects only those transactions with a nonzero system or user completion code.

D|DAY=`nn|nn-nn`

Where `nn` represents a day or range of days as a relative number from the date of the first trace entry (first date is 1, up to a maximum of 99) in the trace buffer or trace log data set. It selects a subset of the entries by a day or range of days.

If there is only one day's data, the END TIME timestamp is `hh:mm:ss.th`. If the trace buffer or trace log data set contains more than one day's data, then END TIME changes to `dd-hh:mm:ss`, where `dd` is the day number relative to the date of the first record in the trace buffer or trace log data set.

If DAY is not specified, but TIME is, the default is the day of the most current record in the buffer or trace log data set.

DB2C|DB2CNT<nnnnnn|>nnnnnn|=nnnnnn

Where nnnnnn represents number of DB2 database calls. It selects a subset of the trace entries in the online trace storage buffer or trace log data set by DB2 database call counts.

DB2T|DB2TIME<nnnnnn|>nnnnnn| =nnnnnn

Where nnnnnn represents DB2 CPU time in milliseconds. It selects a subset of the trace entries in the online trace storage buffer or trace log data set by DB2 elapsed time.

DLIC|DLICNT<nnnnnn|>nnnnnn| =nnnnnn

Where nnnnnn represents number of DL/I database calls. It selects a subset of the trace entries in the online trace storage buffer or trace log data set by DL/I database call counts.

DLIT|DLITIME<nnnnnn|>nnnnnn| =nnnnnn

Where nnnnnn represents DL/I CPU time in milliseconds. It selects a subset of the trace entries in the online trace storage buffer or trace log data set by DL/I elapsed time.

EL|ELAPSED<nnn|>nnn|=nnn

Where nnn represents elapsed time in milliseconds. It selects a subset of the trace entries in the online trace storage buffer or trace log data set by elapsed time.

IOC|IOCNT<nnnnnn|>nnnnnn| =nnnnnn

Where nnnnnn represents number of total I/Os (READs and WRITEs). It selects a subset of the trace entries in the online trace storage buffer or trace log data set by total I/O count.

MSW

Selects only those transactions that are a result of a message switch.

PR|PROGRAM=xxxxxxxx

Where xxxxxxxx represents a program name. It selects a subset of the trace entries in the online trace storage buffer or trace log data set by program name. Multiple trace entries can be selected by using an * or + qualifier; for example, `PR=SM++0+00`.

RE|RESPONSE<nnnnnn|>nnnnnn| =nnnnnn

Where `nnnnnn` represents response time in milliseconds. It selects a subset of the trace entries in the online trace storage buffer or trace log data set by response time.

T|TIME=`hhmm|hhmm-hhmm`

Where `hhmm` represents a start time or time period. It selects a subset of the trace entries in the online trace buffer or trace log data set by a start time or time period.

When the start time is higher than the end time, for example, `TIME=2000-0300`, 24 hours are added.

TR|TRANCODE=`xxxxxxxx`

Where `xxxxxxxx` represents a transaction code. It selects a subset of the trace entries in the online trace storage buffer or trace log data set by transaction name. Multiple trace entries can be selected by using an * or + qualifier; for example, `TR=+MYZ200XX`.

TY|TYPE=`xxx`

Where `xxx` represents a region type which can be:

BMP	Batch message processing and Java batch message processing (JBP) regions.
CON	Conversational (MPP or JMP) regions.
DBT	DBCTL threads (CICS and ODBA).
FPU	Fast Path utility regions.
MDP	Message-driven Fast Path regions.
MPP	Message processing and Java message processing (JMP) regions.
NDP	Non-message-driven Fast Path regions.
TPI	Message processing regions currently executing an explicit CPI-C program.

It selects a subset of the trace entries in the online trace storage buffer or trace log data set by specified region type. Multiple trace entries can be selected by using an * or + qualifier; for example, `TY=M*`.

US|USER=`xxxxxxxx`

Where `xxxxxxxx` represents a user or LTERM ID. It selects a subset of the trace entries in the online trace storage buffer or trace log data set by user or LTERM. Multiple trace entries can be selected by using an * or + qualifier; for example, `US=CICS*`.

Expand: The LTRAC display can be EXPANDed by using the following fields as indicated:

MON(WKLD)
: Active Timer Requests display of all active monitors in the IMS workload (WKLD) area. One of the following fields, HISTORY or CURRENT, appears in the EXPAND line only if the trace request being displayed by LTRAC specified trace logging.

HISTORY
: This field is shown in the EXPAND line of an LTRAC display when the trace request selected from Current Traces (Primary Menu Option 3) specified logging. You can use it to toggle from viewing data from the online storage buffers to viewing entries in the trace log data set for that trace request. It provides access to data no longer in the online trace storage buffer. If the entries are from the online trace storage buffer, line 4 of the display shows `ENTRIES IN BUFFER`.

CURRENT
: This field is shown in the EXPAND line of an LTRAC display showing trace entries from a trace log data set. You can use it to toggle from viewing a trace log data set to viewing any entries remaining in the online trace storage buffer (trace is still active or completed but not purged). If the entries are from a trace log data set, line 4 of the display shows `ENTRIES IN DATASET`. This field applies only to trace requests specifying trace logging.

LINESEL(STRAC)
: You can position the cursor anywhere on a row for a specific transaction occurrence to see a summarized display of trace statistics (STRAC) about that selected occurrence.

Sorting: The display list cannot be sorted.

Scrolling: The display is scrollable up or down. Use PF7 to scroll up to new trace entries as they arrive.

Field Descriptions: Each of the fields is shown and described below in alphabetical order.

DB2
: Total number of DB2 calls issued by this transaction.

DLI
: Total number of DL/I calls issued by this transaction.

ELAPSED
: Elapsed time for this transaction. Time is in milliseconds (ms) unless the value is too high to fit in the field, then it is shown as `hh:mm:ss`.

Note: Time shorter than 10 milliseconds is not available from IMS timestamps and is shown as `<10ms`.

END TIME

Date and time the transaction stopped and the response was queued (synchronization point). The date above the END TIME field is the date of the oldest record that can be displayed. This is either the oldest record in the online trace storage buffer or the oldest record in the current trace log data set.

If the trace record is from the same day as the oldest record, the format of this field is `hh:mm:ss.th`. If the trace record is from a more recent day, the format is `dd-hh:mm:ss`, where `dd` is the day relative to the oldest record.

ENTRIES IN

`ENTRIES IN xxxxxxx n1 - n2` identifies the trace record sequence number in either the online trace storage buffer or the trace log data set, where:

`xxxxxxx` Is BUFFER or DATASET.

`n1` Is the trace record sequence number of the first trace entry.

`n2` Is the trace record sequence number of the last trace entry.

HIST TGT

If the data is being displayed from a trace log data set and the target of the trace in the selected data set is different from the current target shown in the TGT field, this field shows the historical target of the trace. TGT shows the current target controlled by the BBI-SS PAS.

PROGRAM

Program name used by this transaction.

RESPONSE

Response time for this transaction. It is defined as the difference between transaction STOP time and transaction arrival time on the input queue. For a message switch transaction, the arrival time of the original transaction is used, which is indicated by an * following the response time.

Time is in milliseconds (ms) unless the value is too high to fit in the field, then it is shown as `hh:mm:ss`.

Time shorter than 10 milliseconds is not available from IMS timestamps and is shown as `<10ms`

ROW

`ROW n1 OF n2` identifies the number of entries in either the trace log data set or the online trace storage buffer depending upon the trace selected, where:

`n1` Is the entry number of the trace at the top of the display.

`n2` Is the total number of trace entries.

TRANCODE

IMS transaction code

TYP	Type of region processing the transaction:

BMP	Batch message processing region.
CON	Conversational (MPP or JMP) region.
DBT	DBCTL CICS thread.
FPU	Fast Path utility region.
JBP	Java batch message processing region.
JMP	Java message processing region.
MDP	Message-driven Fast Path region.
MPP	Message processing region.
ODB	DBCTL ODBA thread.
NDP	Non-message-driven Fast Path region.
TPI	Message processing region currently executing an explicit CPI-C program.

USR/LTRM	ID of the user who submitted this transaction or LTERM where the transaction was submitted.

Chapter 35. STRAC - Summary Trace Data Display

```
BMC SOFTWARE--------------IMS SUMMARY TRACE--------------------PERFORMANCE MGMT
SERV ==> STRAC           INPUT    11:44:25  INTVL=> 3  LOG=> N  TGT==> IMSxxx
PARM ==> ,SEQ=2                                   LINE 1 OF 100 SCROLL=> CSR
EXPAND:  MON(WKLD), DTRAC, HISTORY                GOTO( DC DLI FP DB2 CPU MSC)
                                             PF10/11 FOR PREV/NEXT TRACE ENTRY
TRANCODE..... THDAMINQ       REGION ID ...        3   HIST TGT......   IMSxxx
PROGRAM...... PHDAMINQ       REGION NAME.. MPPRGN01   USER..........     PWW1
===============================================================================
PSBNAME...... PHDAMINQ       TYPE.........      MPP   LTERM.........  LBS4022
CLASS........        1       ROUTCODE.....            NODENAME......   BS4022
APPC.........       NO       OTMA.........       NO   ABENDCODE....      S0C7
ARRIVAL 15DEC00  09:20:21.01 RESPONSE TIME    3,330ms
START.. 15DEC00  09:20:21.11 ELAPSED TIME.    3,230ms
END ... 15DEC00  09:20:24.34 MSG SWITCH...       NO

----------------------------DC CALL ACTIVITY-----------------------------------
MSG GU........       1      MSG GN........      0      MSG OTHER.....        0
MSG PURGE.....       0      MSG ISRT......      0      MSG TOT.......        1
LAST TRANCODE.  ABMAIL      LAST LTERM....BS234456

----------------------------DL/I CALL ACTIVITY------------  ---I/O----- -------
DBNAME   ORG       GU    GN  REPL  ISRT  DLET OTHER  TOTAL   READ WRITE
-------- ------ ----- ----- ----- ----- ----- -----  -----   ----- -----
DB1H     PHIDAM   280    40   100   100   100    11    631      0     0
 DB1H1X  (DD)                                                  20     0
 DB1H1A  (DD)                                                  40     0
 DB1H2X  (DD)                                                  21     0
 DB1H2A  (DD)                                                  21     0
ALLDBS              0     0     0     0     0     0      0      0    20
**TOTAL           280    40   100   100   100    11    631    102    20

----------------------------FAST-PATH ACTIVITY---------------------------------
NBA..........       10      CI CONTENTION       0      SYNC RETCODE..        0
OBA..........        2      WAIT FOR BFRS       0      BALG Q-COUNT          0
BUFFER USED..        0      WAIT-OBA BFR LATCH  NO

----------------------------DB2 ACTIVITY---------------------------------------
DB2NAME...... DB2D          CONTROL.......      0      INSERTS.......        0
PLANNAME..... PHDAMINQ      DYNAMIC.......      0      DELETES.......        0
SEL/FCH......        5      DDL...........      0      UPDATES.......        0
OPENS........        0      OTHER.........      0      SQL TOTAL.....        5

----------------------------CPU TIMES------------------------------------------
                   DLI TIME   BUFFER HANDLER TIME   OPEN CLOSE TIME    TOTAL
DEPENDENT REGION     200us              0                 0          200us
CONTROL REGION       400us              0                 0          400us
DLISAS               300us              0                 0          300us
DEPENDENT REGION CPU TIME..........................................  100ms
DEPENDENT REGION DB2 CPU TIME......................................   50ms
TOTAL CPU TIME.....................................................  151ms

----------------------------MSC ACTIVITY---------------------------------------
ARRIVAL DATE..15DEC00       TIME......... 09:20:21      INPUT SYSID...       1
DESTINATION SYSID:
IOPCB.........       2      MSG SWITCH....      1      ALTERNATE PCB.        2
***************************** END OF DATA **************************************
```

Description: This service summarizes DC call activity, DL/I call activity, Fast Path activity, DB2 activity, and CPU times for a selected transaction instance. You can use it, for example, to see how many times this transaction instance accessed a database, the types of calls it made and how many, and the transaction elapsed time. An unusually high elapsed transaction time may indicate poor DASD response time.

Note: In the CPU timing area of the display, to see the amount of dependent, control, and DLISAS region CPU time used by the:

- DL/I analyzer to process a transaction and open and close databases
- Buffer handler module to process a transaction

you must set the Event Collector options in BBPARM member IMFECP00 and the IMS LSO option as described for these fields on pages 283 and 283. For information about these options, see the *MAINVIEW for IMS Online – Customization Guide*.

You can request STRAC by cursor selection from the LTRAC display or from the analyzer service list (Primary Menu Option 1).

Color: If you have a color monitor:

Turquoise	Indicates normal values.
Red	Highlights any transactions that have abended.
White	Indicates column headings.
Yellow	Indicates error messages.

Select Code: STRAC

Parameter: The STRAC parameters can be used as follows:

- Multiple parameters must be separated by commas.
- A blank indicates the end of a parameter string.
- Any filter parameter used for LTRAC (see the parameter descriptions in Chapter 34, “LTRAC - List of Trace Entries” on page 265) also can be used for STRAC.

traceid

Where traceid is a unique ID used with the MTRAC request. A blank can be an identifier. This parameter is positional and must be in the first position. This field is primed automatically with the MTRAC request identifier if the request is selected from the Current Traces or Active Timer requests application.

SEQ=n

Where n is a unique internal sequence number assigned to each trace entry in the online trace storage buffer or trace log data set. The number is assigned sequentially throughout the trace. If logging is requested for the trace, the number is assigned across multiple log data sets. If this keyword is not entered, the first trace entry is displayed. If you request STRAC by EXPANDing from LTRAC, this field is primed automatically. You can change the value to scroll anywhere through the trace data.

Expand: The STRAC display can be EXPANDed to the following displays by cursor selection:

MON(WKLD)
: Active Timer Requests display of all active monitors in the IMS workload (WKLD) area.

DTRAC
: Detail trace entry display (DTRAC service).

HISTORY
: This field is shown in the EXPAND line of an STRAC display if the selected trace request specified logging. You can use it to toggle from viewing data from the online storage buffers to viewing entries in the trace log data set. It provides access to data no longer in the online trace storage buffer.

CURRENT
: This field is shown in the EXPAND line of an STRAC display showing trace entries from a trace log data set. You can use it to toggle from viewing a trace log data set to viewing any entries remaining in the online trace storage buffer (trace is still active or completed but not purged). This field applies only to trace requests specifying trace logging.

GOTO: You can go to a specific area of the display by moving it to the top. Move the cursor to any of the following GOTO fields in the EXPAND line and press ENTER to move that area to the top of the display:

DC	DC ACTIVITY (This field does not apply to an IMS DBCTL system.)
DLI	DL/I CALL ACTIVITY
FP	FAST-PATH ACTIVITY
DB2	DB2 ACTIVITY
CPU	CPU TIMES
MSC	MSC ACTIVITY (This field does not apply to an IMS DBCTL system.)

Sorting: This display cannot be sorted.

Scrolling: You can scroll the display by:

- Using PF keys to scroll up or down.

Using PF10 and PF11 to scroll through trace entries, meeting specified filter parameter criteria, in the online trace storage buffer or trace log data set. PF10 scrolls to the previous trace entry. PF11 scrolls to the next trace entry. The current entry is displayed in the SEQ field.

- Selecting a GOTO field which selects a display area and scrolls it to the top as described in "GOTO:".

Field Descriptions: Each of the fields is shown and described below by display area.

Area 1

```
TRANCODE..... THDAMINQ       REGION ID ...         3    HIST TGT......   IMSxxx
PROGRAM...... PHDAMINQ       REGION NAME.. MPPRGN01     USER..........     PWW1
===============================================================================
PSBNAME...... PHDAMINQ       TYPE.........       MPP    LTERM.........  LBS4022
CLASS........        1       ROUTCODE.....              NODENAME......   BS4022
APPC.........       NO       OTMA.........        NO    ABENDCODE....      S0C7
ARRIVAL 15DEC00  09:20:21.01 RESPONSE TIME    3,330ms
START.. 15DEC00  09:20:21.11 ELAPSED TIME.    3,230ms
END ... 15DEC00  09:20:24.34 MSG SWITCH...        NO
```

This area shows general information about the application program activity in the region for the selected transaction instance. Data above the ==== line cannot be scrolled.

The descriptions are arranged in alphabetical order.

ABENDCODE
System abend code (Sxxx) or user abend code (Uxxxx) if the transaction terminated abnormally.

APPC
YES|NO. Indicates whether the original message was received by an LU6.2 or APPC connection.

ARRIVAL
Date and time of transaction arrival. For a message switch, the original transaction arrival time is used. This field is blank for Fast Path transactions.

CLASS
Class used by this transaction. This is always 0 for a DBCTL system.

ELAPSED TIME
Current time minus the transaction start time. It is in milliseconds (ms) or hh:mm:ss if the millisecond value is too large to fit in the field.

Note: Time shorter than 10 milliseconds is not available from IMS timestamps and is shown as <10ms.

END
: Date and time the transaction stopped and the response was queued (synchronization point).

HIST TGT
: If the request specified trace logging, this field shows the historical target from the trace log data set if it is different from the current target.

LTERM
: LTERM name of the terminal where this transaction was submitted.

  The name of the LTERM that originated the transaction.

  For DBCTL threads, the LTERM is set to N/A.

  For a non-message-driven BMP or JBP, the LTERM is set to *NONE*.

MSG SWITCH
: YES|NO indicates whether this transaction is a result of a message switch.

NODENAME
: VTAM node of the originating terminal for this transaction.

OTMA
: YES|NO indicates whether the type of CNT for this LTERM is Open Transaction Manager Access.

PROGRAM
: Name of the application program called by IMS to process this transaction.

PSBNAME
: PSB name of the application program used to process this transaction.

REGION ID
: Region identifier assigned to this region by IMS.

REGION NAME
: Job or STC name of the message region where the transaction executed.

RESPONSE TIME
: Response time in milliseconds (ms) or hh:mm:ss if the millisecond value is too large to fit in the field.

  It is calculated by subtracting the message arrival time from the transaction stop time.

  Time shorter than 10 milliseconds is not available from IMS timestamps and is shown as <10ms.

ROUTCODE
: Fast path routing code if applicable.

START
: Date and time the transaction started.

TRANCODE

Name of this transaction. For a non-message-driven BMP or JBP, the transaction code is set to blank.

TYPE

Type of region processing the transaction:

BMP	Batch message processing region.
CON	Conversational MPP or JMP region.
DBT	DBCTL CICS thread.
FPU	Fast Path utility regions.
JBP	Java batch message processing region.
JMP	Java message processing region.
MDP	Message-driven Fast Path region.
MPP	Message processing region.
NDP	Non-message-driven Fast Path region.
ODB	DBCTL ODBA thread.
TPI	Message processing region currently executing an explicit CPI-C program.

USER

ID of the user who requested this transaction.

Area 2

```
----------------------------DC ACTIVITY----------------------------------------
MSG GU........       1      MSG GN........       0      MSG OTHER.....      0
MSG PURGE.....       0      MSG ISRT......       0      MSG TOT.......      1
LAST TRANCODE.  ABMAIL      LAST LTERM.... BS234456
```

This area shows the amount of DC call activity incurred by this transaction instance. The descriptions are arranged in alphabetical order.

LAST LTERM

Name of the last LTERM to which a message was sent by this transaction instance.

LAST TRANCODE

Transaction code of the last message switch transaction inserted by this transaction instance, if any.

MSG GN

Number of MESSAGE GET NEXT calls issued by this transaction instance. For MDPs, this value is 0.

MSG GU

Number of MESSAGE GET UNIQUE calls issued by this transaction. instance. For MDPs (message-driven program), this value is 0.

MSG ISRT

Number of MESSAGE INSERT calls issued by this transaction instance. For MDPs, this value does not include I/O PCB ISRTs.

MSG OTHER

The number of other types of message calls, such as checkpoint, statistics, or system service, issued by this transaction instance.

MSG PURGE

Number of MESSAGE PURGE calls issued by this transaction instance.

MSG TOT

Total number of message calls issued by this transaction instance

.

Area 3

```
----------------------------DL/I CALL ACTIVITY------------  ---I/O----- -------
DBNAME   ORG       GU    GN  REPL  ISRT  DLET OTHER  TOTAL   READ WRITE
-------- ------ ----- ----- ----- ----- ----- -----  -----  ----- -----
DB1H     PHIDAM   280    40   100   100   100    11    631      0     0
 DB1H1X  (DD)                                                  20     0
 DB1H1A  (DD)                                                  40     0
 DB1H2X  (DD)                                                  21     0
 DB1H2A  (DD)                                                  21     0
ALLDBS              0     0     0     0     0     0      0      0    20
**TOTAL           280    40   100   100   100    11    631    102    20
```

This area shows the amount of DL/I call activity incurred by this transaction instance. The descriptions are arranged in alphabetical order.

DBNAME

Name of the DL/I database accessed by this transaction instance.

DLET

Number of DELETE calls issued by this transaction instance against the database.

GN

Number of GET NEXT calls issued by this transaction instance against the database.

GU

Number of GET UNIQUE calls issued by this transaction instance against the database.

ISRT

Number of INSERT calls issued by this transaction instance against the database.

ORG

IMS database organization type. If DBTNAME=DD is specified in BBPARM member IMFECP00, ddnames are shown as (DD); for example, CUSTHISM (DD). If I/O was done against the database or data set, the type and amount of I/O is shown by the READ or WRITE fields.

DEDB	DEDB
HD-R	HDAM plus root-index
HDAM	HDAM

HID-I	HIDAM index
HID-R	HIDAM plus root-index
HIDAM	HIDAM
HIS-C	HISAM-case-2
HISAM	HISAM
HSAM	HSAM
MSDBD	MSDB, related, dynamic
MSDBF	MSDB, related, fixed
MSDBK	MSDB, nonrelated, key in segment
MSDBL	MSDB, nonrelated, LTERM key
PHDAM	Partitioned HDAM
PHIDAM	Partitioned HIDAM
PSINDX	Partitioned secondary index
ROOT	Root-index
SHISAM	SHISAM
SSAM	SSAM

OTHER
: Number of other types of DL/I calls issued against this database by this transaction instance.

READ
: Total number of key reads and nonkey reads to satisfy DL/I requests for this database. STRAC reports only application-related I/O; DTRAC reports all I/O including I/O done by IMS at database OPEN.

REPL
: Number of REPLACE calls issued by this transaction instance against the database. This row contains totals of the requests for all databases accessed by this transaction instance.

TOTAL
: Total number of DL/I database calls issued by this transaction instance against the database:

GU + GN + REPL + ISRT + DLET

This value includes calls to MSDBs and DEDBs. For large values, the number is scaled by 1000 and a K is placed after the value.

WRITE
: Total number of key writes and nonkey writes to satisfy DL/I requests for this database. STRAC reports only application-related I/O; DTRAC reports all I/O including I/O done by IMS at database OPEN.

Area 4

```
 ----------------------------FAST-PATH ACTIVITY---------------------------------
NBA..........       10      CI CONTENTION.       0      SYNC RETCODE..   0
OBA..........        2      WAIT FOR BFRS.       0      BALG Q-COUNT..   0
BUFFER USED..        0      WAIT-OBA BFR LATCH. NO
```

This area shows the region Fast Path call activity for the selected transaction instance. The descriptions are arranged in alphabetical order.

BALG Q-COUNT

For a Fast Path message-driven transaction, the number of transactions queued on the same balancing group when this transaction went through synchronization point processing.

BUFFER USED

Number of Fast Path buffers actually used by this transaction instance.

Tuning Tip

If this total is greater than the NBA value, the extra buffers are taken from the overflow buffer allocation. The use of overflow buffers should be avoided if possible.

CI CONTENTION

Number of Fast Path control interval contentions for this transaction instance.

NBA

Number of buffers defined in the region JCL for normal buffer allocation.

OBA

Number of buffers defined in the region JCL for overflow buffer allocation.

SYNC RETCODE

The result of synchronization point processing for a transaction that accessed Fast Path databases.

WAIT FOR BFRS

Number of times this transaction instance had to wait for a Fast Path buffer to become available.

WAIT-OBA BFR LATCH

YES|NO indicates whether this transaction waited for the OBA buffer latch or not.

Tuning Tip

Only one message region at a time can use the overflow buffers. A Y in this field indicates that this region used all of the normal buffers allocated to it (NBA) and had to wait before getting one or more overflow buffers. If this wait is a common occurrence, you should consider increasing the NBA for this region so that overflow buffers are not needed.

Area 5

```
 ----------------------------DB2 ACTIVITY----------------------------------------
DB2NAME...... DB2D          CONTROL.......       0      INSERTS.......       0
PLANNAME..... PHDAMINQ      DYNAMIC.......       0      DELETES.......       0
SEL/FCH......        5      DDL...........       0      UPDATES.......       0
OPENS........        0      OTHER.........       0      SQL TOTAL.....       5
```

This area shows the region DB2 subsystem call activity for the selected transaction instance. For DBCTL, only BMPs are shown. CICS does not use IMS services to access DB2.
The descriptions are arranged in alphabetical order.

CONTROL
: Total number of SQL control type calls issued to DB2 (GRANT or REVOKE, for example) by this transaction instance.

DB2NAME
: DB2 subsystem ID.

DDL
: Number of SQL Data Definition Language calls issued to DB2 (CREATE, DROP, ALTER, COMMENT, or LABEL, for example) by this transaction instance.

DELETES
: Total number of SQL DELETE calls issued to DB2 by this transaction instance.

DYNAMIC
: Total number of SQL dynamic calls issued to DB2 (PREPARE, DESCRIBE, or EXECUTE, for example) by this transaction instance.

INSERTS
: Total number of SQL INSERT calls issued to DB2 by this transaction instance.

OPENS
: Total number of SQL OPEN cursor calls issued to DB2 by this transaction instance.

OTHER
: Number of other SQL calls issued to DB2 by this transaction instance that do not fit in any of the other types defined in this display (EXPLAIN, LOCK, LABEL, CLOSE, or table and security manipulation, for example).

PLANNAME
: Plan name used for this transaction.

SEL/FCH
: Total number of SQL SELECT and FETCH calls issued to DB2 by this transaction instance.

SQL TOTAL
: Total number of SQL calls issued to DB2 by this transaction instance. For large values, the number is scaled by 1000 and a K is placed after the value.

UPDATES
: Total number of SQL UPDATE calls issued to DB2 by this transaction instance.

```
 Area 6

 ----------------------------CPU TIMES-----------------------------------------
                   DLI TIME   BUFFER HANDLER TIME   OPEN CLOSE TIME    TOTAL
DEPENDENT REGION     200us              0                 0           200us
CONTROL REGION       400us              0                 0           400us
DLISAS               300us              0                 0           300us
DEPENDENT REGION CPU TIME.........................................    100ms
DEPENDENT REGION DB2 CPU TIME.....................................     50ms
TOTAL CPU TIME....................................................    151ms
```

This area shows how much CPU time is used by the DL/I Analyzer, the buffer handler module, and the application program to process a transaction.

The top portion of the area:

```
                     DLI TIME   BUFFER HANDLER TIME   OPEN CLOSE TIME    TOTAL
DEPENDENT REGION      200us              0                  0            200us
CONTROL REGION        400us              0                  0            400us
DLISAS                300us              0                  0            300us
```

shows the DL/I Analyzer, buffer handler module, and application program usage of CPU time. These fields, in alphabetical order, provide:

CONTROL REGION - BUFFER HANDLER TIME
: Amount of control region CPU time used by the buffer handler module to process this transaction. To obtain this time, the Event Collector BHTO option in BBPARM member IMFECP00 must be set to:

  BHTO=ON

  If it is not, 0s are displayed for this field.

CONTROL REGION - DLI TIME
: Amount of control region CPU time used by the DL/I Analyzer to process this transaction. To obtain this time, the Event Collector CPU option in BBPARM member IMFECP00 must be set to:

  CPU=ALL

  If it is not, 0s are displayed for this field.

CONTROL REGION - OPEN CLOSE TIME
: Amount of control region CPU time used by the DL/I Analyzer to open and close databases for this transaction. To obtain this time, the Event Collector CPU option in BBPARM member IMFECP00 must be set to:

  CPU=ALL

  and the IMS LSO option must be set to:

  LSO=YES

  If they are not, 0s are displayed for this field.

DEPENDENT REGION - BUFFER HANDLER TIME
: Amount of dependent region CPU time used by the buffer handler module to process this transaction. To obtain this time, the Event Collector BHTO option in BBPARM member IMFECP00 must be set to:

  BHTO=ON

  If it is not, 0s are displayed for this field.

DEPENDENT REGION - DLI TIME

Amount of dependent region CPU time used by the DL/I Analyzer to process this transaction. To obtain this time, the Event Collector CPU option in BBPARM member IMFECP00 must be set to:

```
CPU=ALL
```

or

```
CPU=DEP
```

If it is not, 0s are displayed for this field.

DEPENDENT REGION - OPEN CLOSE TIME

Amount of dependent region CPU time used by the DL/I Analyzer to open and close Fast Path databases for this transaction.

DLISAS REGION - BUFFER HANDLER TIME

Amount of DLISAS region CPU time used by the buffer handler module to process this transaction. To obtain this time, the Event Collector BHTO option in BBPARM member IMFECP00 must be set to:

```
BHTO=ON
```

If it is not, 0s are displayed for this field.

DLISAS REGION - DLI TIME

Amount of DLISAS region CPU time used by the DL/I Analyzer to process the transaction. To obtain this time, the Event Collector CPU option in BBPARM member IMFECP00 must be set to:

```
CPU=ALL
```

If it is not, 0s are displayed for this field.

DLISAS REGION - OPEN CLOSE TIME

Amount of DLISAS region CPU time used by the DL/I Analyzer to open and close databases for this transaction. To obtain this time, the Event Collector CPU option in BBPARM member IMFECP00 must be set to:

```
CPU=ALL
```

If it is not, 0s are displayed for this field.

TOTAL

The sum of the DL/I, buffer handler, and open close CPU time used by the dependent, control, or DL/I SAS regions to process a transaction, where us represents microseconds, ms represents milliseconds, and s represents seconds.

The bottom portion of this area:

```
DEPENDENT REGION CPU TIME.........................................  100ms
DEPENDENT REGION DB2 CPU TIME.....................................   50ms
TOTAL CPU TIME....................................................  151ms
```

shows application program and DB2 call usage of CPU time. These fields, in alphabetical order, provide:

DEPENDENT REGION CPU TIME
: Amount of CPU time used by the application program in the dependent region to process this transaction. It includes DB2 CPU time.

DEPENDENT REGION DB2 CPU TIME
: Amount of CPU time used in the dependent region to process normal and service calls to DB2.

TOTAL CPU TIME
: The sum of all CPU time used.

Area 7

```
----------------------------MSC ACTIVITY----------------------------------------
ARRIVAL DATE..15DEC00       TIME......... 09:20:21      INPUT SYSID...        1
DESTINATION SYSID:
IOPCB.........       2      MSG SWITCH....       1      ALTERNATE PCB.        2
```

This area shows system information from the processing of a multiple systems coupling (MSC) transaction. These fields, in alphabetical order, provide:

ARRIVAL DATE - TIME

Date and time this MSC transaction was submitted for processing on the originating system.

DESTINATION SYSID:

ALTERNATE PCB
: System ID for the last message sent to a remote LTERM using an alternate I/O PCB.

IOPCB
: ID of the system where output is sent through the I/O PCB.

MSG SWITCH
: ID of the system where the last message switch transaction was inserted.

INPUT SYSID
: ID of the system where this transaction was entered.

Chapter 36. DTRAC - Detail Trace Data Display

```
BMC SOFTWARE  -------------    IMS DETAIL TRACE     -----------PERFORMANCE MGMT
SERV ==> DTRAC            INPUT   14:24:03  INTVL=> 3  LOG=> N  TGT==> IMSxxx
PARM ==> TRAN1,SEQ=2,LEVEL=1                          ROW 1 OF 6 SCROLL=> CSR
EXPAND:  DLI, IO, DATA, ALL
                                          (PF10/11 FOR PREV/NEXT TRACE ENTRY)
TRANCODE: THDAMINQ  USR/LTRM: PDRIVER   ARRIVE: 14:19:09.00  START: 14:19:09.05
     EVENT              AT      ELAPSED       CPU     DETAIL
------------------ --------- ----------- ----------- -------------------------
SCHD RGN IMSM71X         0us       208ms     1,758ms IWAITING
GU   DC            36,290us        20us         8us ok
GU   DB  CUSTHDAM  36,318us    80,464us        71us ok
GNP  DB  CUSTHDAM      114ms        69us        47us ok
ISRT DC                114ms        24us        21us QH
GU   DC  I/O PCB       114ms     2,570us         0us
```

```
BMC SOFTWARE  -------------    IMS DETAIL TRACE     -----------PERFORMANCE MGMT
SERV ==> DTRAC            INPUT   14:24:44  INTVL=> 3  LOG=> N  TGT==> IMSxxx
PARM ==> TRAN1,SEQ=2,LEVEL=2                          ROW 1 OF 8 SCROLL=> CSR
EXPAND:  DLI, IO, DATA, ALL                                 HIST TGT==>IMSxxx
                                          (PF10/11 FOR PREV/NEXT TRACE ENTRY)
TRANCODE: THDAMINQ  USR/LTRM: PDRIVER   ARRIVE: 14:19:09.00  START: 14:19:09.05
     EVENT              AT      ELAPSED       CPU     DETAIL
------------------ --------- ----------- ----------- -------------------------
SCHD RGN IMSM71X         0us       208ms     1,758ms IWAITING
GU   DC            36,290us        20us         8us ok
GU   DB  CUSTHDAM  36,318us    80,464us        71us ok
  IO DD  CUSTHDAM      102ms     7,669us
  IO DD  CUSTHDAM      110ms     3,696us
GNP  DB  CUSTHDAM      114ms        69us        47us ok
ISRT DC                114ms        24us        21us QH
GU   DC  I/O PCB       114ms     2,570us         0us
```

```
BMC SOFTWARE  -------------    IMS DETAIL TRACE     -----------PERFORMANCE MGMT
SERV ==> DTRAC            INPUT   14:25:12  INTVL=> 3  LOG=> N  TGT==> IMSxxx
PARM ==> TRAN1,SEQ=2,LEVEL=3                         ROW 1 OF 29 SCROLL=> CSR
EXPAND:  DLI, IO, DATA, ALL
                                          (PF10/11 FOR PREV/NEXT TRACE ENTRY)
TRANCODE: THDAMINQ  USR/LTRM: PDRIVER   ARRIVE: 14:19:09.00  START: 14:19:09.05
     EVENT              AT      ELAPSED       CPU     DETAIL
------------------ --------- ----------- ----------- -------------------------
SCHD RGN IMSM71X         0us       208ms     1,758ms IWAITING
GU   DC            36,290us        20us         8us ok
     IOA  m  THDAMINQ ARMCO STEEL
         0900ECCCDCDD4CDDCD4EECCD444444444444444444444444444444444444444444
         0430384149580194360235530000000000000000000000000000000000000000
GU   DB  CUSTHDAM  36,318us    80,464us        71us ok
     SSA CUSTOMER(CUSTID   =ARMCO STEEL              )
         CEEEDDCD4CEEECC4447CDDCD4EECCD444444444444445
         34236459D342394000E194360235530000000000000D
     KFB ARMCO STEEL
         CDDCD4EECCD4444444444444
         1943602355300000000000000
     IOA ARMCO STEEL
         CDDCD4EECCD444444444444444444400000000000000000000000000000000000
         194360235530000000000000000000000000000000000000000000000000000000
GNP  DB  CUSTHDAM      114ms        69us        47us ok
     IOA CIMS      121075    17000
         CCDE444444FFFFFF4444FFFFF444440000000000000000000000000000000000
         3942000000121075000017000000000000000000000000000000000000000000
ISRT DC                114ms        24us        21us QH
GU   DC  I/O PCB       114ms     2,570us         0us
```

```
BMC SOFTWARE  -------------    IMS DETAIL TRACE     -----------PERFORMANCE MGMT
SERV ==> DTRAC           INPUT   14:25:12  INTVL=> 3  LOG=> N  TGT==> IMSxxx
PARM ==> TRAN2,SEQ=2,LEVEL=2                         ROW 1 OF 29 SCROLL=> CSR
EXPAND:  DLI, IO, DATA, ALL
                                           (PF10/11 FOR PREV/NEXT TRACE ENTRY)
TRANCODE: CSQDB2T   USR/LTRM: WTOR      ARRIVE: 13:37:09.90  START: 13:37:21.91
     EVENT             AT      ELAPSED       CPU     DETAIL
------------------ --------- ----------- ----------- -------------------------
SCHD RGN IMSM19X         0us     6,865us         0us IWAITING
GU   DC  WTOR       25,314us        25us         8us ok
CSQ3     SIGN ON    48,334us     2,535us       269us RC: 0
CSQ3     C_THREAD   51,057us        17us        12us RC: 0
CSQ3     MQCONN     51,109us        37us        33us RC: 0
CSQ3     MQPUT1     51,367us     2,645us       712us RC: 0
DB2S     SIGN ON    54,195us       129us        81us RC: 0
DB2S     C_THREAD   54,368us       294ms     1,698us RC: 0
DB2S     OPEN          347ms       578us       338us RC: 0   SQL-307
DB2S     FETCH         348ms       161ms     1,139us RC: 0   SQL-317
DB2S     FETCH         509ms       864us       135us RC: 0   SQL-317
DB2S     FETCH         510ms        73us        70us RC: 0   SQL-317
DB2S     FETCH         510ms       278us        62us RC: 0   SQL-317
DB2S     FETCH         512ms       149us       103us RC: 100 SQL-317
DB2S     CLOSE         512ms        70us        49us RC: 0   SQL-327
CSQ3     MQDISC        542ms       136us       132us RC: 0
CSQ3     P1-COMIT      544ms        95us        91us RC: 0
DB2S     P1-COMIT      544ms       312us       218us RC: 0
CSQ3     T_THREAD      545ms     5,207us       177us RC: 0
DB2S     T_THREAD      550ms       605us       245us RC: 0
```

Description: This service displays a detail trace of events that occurred as a transaction was processed. It shows data for IMS events and calls to IMS by DB2 and MQSeries.

The detail workload trace is used to debug complex IMS performance problems associated with the execution of specific transactions. Unique events in the flow of each traced transaction are displayed.

Usage notes: The detail trace data is captured during execution into detail trace buffers allocated when the trace is started. The TRSIZE and TRBUFF parameters in BBPARM member IMFBEX00 control the size and number of the buffers.

The detail trace buffer is acquired for a transaction either during scheduling or when the transaction first starts running. If the trace did not complete initialization or if a buffer is not available when the transaction begins running, no detail trace entries will be recorded for the transaction. When you try to view the detail trace, you will get the following message:

```
WA3141I DETAIL TRACE NOT ACTIVE FOR THIS TRANSACTION
        EXECUTION
```

The number of detail trace entries that will fit into the detail trace buffer depends on

- the size of the buffer allocated
- the size of the entries captured

The size of the entry will vary based on its type. If you specified WRAP=NO when you started the trace (see page 208) and a transaction runs out of space in the buffer, you will get the following message:

```
WA3143I DETAIL TRACE BUFFER FULL
```

Select Code: DTRAC

Parameters: The DTRAC parameters can be used as follows:

- Multiple parameters must be separated by commas.
- A blank indicates the end of a parameter string.

traceid

The traceid is a unique ID used with the MTRAC request. A blank can be an identifier. This parameter is positional and must be in the first position. This field is primed automatically with the MTRAC request identifier if DTRAC is requested from the STRAC display.

LEVEL=1|2|3|4

Level 1 displays DL/I call activity for transaction events. Level 2 (the default) includes DL/I call activity and, for IMS calls, displays database I/O events that occurred during call processing. Level 3 includes DL/I call activity and, for IMS calls, displays segment search argument, key feedback, and I/O area data (SSA, KFB, and IOA). Level 4 displays all the information of Levels 1, 2, and 3.

Levels 1, 2, 3, and 4 correspond with the DLI, IO, DATA, and ALL options on the EXPAND line.

SEQ=n

The SEQ= value is a unique internal sequence number assigned sequentially to each trace entry in the online trace storage buffer or trace log data set. If logging was requested for the trace, the number is assigned across multiple log data sets. If this keyword is not entered, the most recent trace entry is displayed. If you request DTRAC by navigating from STRAC, this field is primed automatically. You can change the value to scroll anywhere through the trace data.

Expand: You can select the following DTRAC display options by tabbing on the EXPAND line:

DLI — This expansion option displays only the DL/I call activity for transactions.

IO — This is the default expansion option. It includes DL/I call activity, and for IMS transactions, it displays database I/O events that occurred during call processing for the transactions.

DATA — This expansion option includes DL/I call activity, and for IMS transactions, it displays segment search argument, key feedback, and I/O area data (SSA, KFB, and IOA).

ALL — This expansion option includes DL/I call activity, and for IMS transactions, it displays database I/O events and SSA, KFB, and IOA data.

The DLI, IO, DATA, and ALL options on the EXPAND line correspond with Levels 1, 2, 3, and 4 on the PARM line (see the LEVEL= parameter on page 289).

HISTORY is displayed in the EXPAND line if the selected trace request specified logging. You can use it to toggle from viewing data from the current trace storage buffer to viewing entries in the trace log data set. It provides access to data no longer in the current trace storage buffer (see Chapter 28, “Logging a Trace”).

CURRENT is displayed on the EXPAND line when a DTRAC display shows trace entries from a trace log data set. You can use it to toggle from viewing a trace log data set to viewing any entries remaining in the current trace storage buffer (trace is still active or completed but not purged). This field applies only to trace requests specifying trace logging.

Sorting: The DTRAC display cannot be sorted.

Scrolling: There are two ways to scroll DTRAC displays.

- You can use PF keys to scroll up or down.
- You can use PF10 and PF11 to scroll through trace entries in the current trace storage buffer or trace log data set. PF10 scrolls to the previous trace entry. PF11 scrolls to the next trace entry. The current entry is displayed in the SEQ field.

Field Descriptions: The following fields are displayed by DTRAC:

ARRIVE	Time of transaction arrival in the input queue (not applicable to an IMS DBCTL system)
HIST TGT	If the request specified trace logging, this field shows the historical target from the trace log data set if it is different from the current target shown in the TGT field. The current target is controlled by the BBI-SS PAS.
START	Time transaction started
TRANCODE	Name of this transaction instance
USR/LTRM	ID of the user or the name of the LTERM where the transaction was submitted if there is no user ID

Detail Trace Data Columns

This section describes the information DTRAC displays under its five data columns.

```
     EVENT             AT       ELAPSED       CPU      DETAIL
------------------ --------- ----------- ----------- --------------------------
```

EVENT	This field identifies the event being traced. Formatting for information in this column is described on pages 292 through 297.
AT	This field shows the elapsed time from the end of scheduling of the transaction until the time the event occurred. The time value is displayed in seconds (s), milliseconds (ms), or microseconds (us). Certain DB2 events can occur while a region is in an IWAIT status during in scheduling. Those events are traced, but the elapsed time shown is from the start of scheduling. One of the causes of this type of event is when DB2 terminates while a region is signed on to DB2. **Note:** The elapsed time of scheduling events (SCHD) is not considered part of the transaction execution time and is not included in the elapsed time shown in the AT field.
ELAPSED	This field shows the elapsed time of the event in seconds (s), milliseconds (ms), or microseconds (us). **Note:** The elapsed time of a DL/I call includes the elapsed time of any I/Os performed on its behalf. For SCHD entries, elapsed time excludes IWAIT time.
CPU	This field provides the CPU time of an event if the time value is available. The time is displayed in microseconds (us). For scheduling events (SCHD), the CPU field shows the amount of IWAIT (IMS wait) time during scheduling (the amount of time the region waited to schedule a transaction for processing). **Note:** For wait-for-input (SCHD WFI) transactions, the event is at the end of the trace events. Because of the method used to collect this data, the data (trace buffer) is available for display only when the next transaction is scheduled in that region. The trace buffer is transferred to the PAS when the WFI GU to the I/O PCB is satisfied or the region terminates.
DETAIL	For IMS nonscheduling events, this field displays the DL/I status code returned by IMS for the event. For IMS scheduling events (SCHD), IWAITING is displayed in the field. For DB2 events, the detail field displays the return code associated with the event. For SQL DB2 events, the field also displays the DB2 precompiled SQL statement number. For MQSeries events, the field displays the return code associated with the event and the token, if there is one.

Detail Trace Event and Data Display Formats

This section describes the DTRAC event and data display formats for trace information. DTRAC displays detail trace information for:

- IMS calls
- DB2 calls
- CICS calls
- MQSeries calls

IMS Calls – Event Display Formats

This section describes the formatting used in the DTRAC EVENT column for events resulting from IMS calls. The types of IMS calls are:

- data communication calls
- database calls
- scheduling calls

For IMS calls, DTRAC also displays

- database I/O events that occurred during call processing for transactions (page 293)
- segment search argument (SSA) data (page 294)
- key feedback (KFB) data (page 294)
- I/O area (IOA) data (page 294)

Data Communications Calls

Shown below is the format DTRAC uses for IMS data communications calls.

Format:

- Call type (such as GU, GN, ISRT, PURG)
- DC (message queue) or
  EM (expedited message handler)
- Name of the PCB referenced in the call, such as I/O PCB

Examples:

```
ISRT DC I/O PCB
GU   EM I/O PCB
```

Database Calls

Shown below is the format DTRAC uses for IMS database calls.

Format:

- Call type (such as GU, GN, ISRT, DLET, REPL)
- DB (DL/I database),
  DE (Fast Path DEDB), or
  MS (Fast Path MSDB)
- Name of the database (the DBD name) referenced in the calls

Examples:

```
REPL DB CUSTHDAM
GU   DE PAYDEDB
```

Scheduling Calls

Shown below is the format DTRAC uses for IMS scheduling calls.

Format:

- SCHD
- RGN (region) or
  WFI (wait for input)
- Region name of this IMS dependent region

Example:

```
SCHD RGN IIMSM71X
```

IMS Calls – Data Display Formats

Database I/O

DTRAC displays the following IMS database I/O.

Format:

- IO DD
- ddname of the file for which I/O was performed

Note: Fast Path DEDB I/O is not traced.

Example:

```
IO DD CUSTDD
```

Segment Search Arguments (SSA)

DTRAC displays the following IMS SSA data.

Format:

- SSA
- First 45 bytes of the nth SSA, in hexadecimal dump format, specified with the call. All SSAs specified for a call are displayed.

Example:

```
SSA CUSTOMER(CUSTID   =ARMCO STEEL              )
    CEEEDDCD4CEEECC4447CDDCD4EECCD444444444444445
    34236459D342394000E19436023553000000000000000D
```

Key Feedback Data

DTRAC displays the following IMS key feedback data.

Format:

- KFB
- First 25 bytes in hexadecimal dump format of the key feedback data specified with the call

Example:

```
KFB ARMCO STEEL
    CDDCD4EECCD44444444444444
    1943602355300000000000000
```

I/O Area (IOA)

DTRAC displays the following IMS I/O area data.

Format:

- IOA
- First 65 bytes in hexadecimal dump format of the I/O area specified with the call

Example:

```
IOA ARMCO STEEL
    CDDCD4EECCD4444444444444444444400000000000000000000000000000000000
    19436023553000000000000000000000000000000000000000000000000000000
```

DB2 Calls

This section describes the formatting used in the DTRAC EVENT column for events resulting from DB2 calls. The types of DB2 calls are:

- IMS attach facility DB2 calls
- SQL DB2 calls

IMS Attach Facility DB2 Calls

Shown below is the format DTRAC uses for service calls issued by the IMS attach facility to a DB2 subsystem.

Format:

- DB2 subsystem identifier
- IMS attach facility call type

Example:

```
DB2S  C_THREAD
```

IMS attach facility call types:

`xx-UNKC`	Unrecognized call type; `xx` is the hex value of the call type.
`ABORT`	Abort call
`C_THREAD`	Create thread
`IDENTIFY`	Identify
`INIT_CAL`	Initiate call
`P1-COMIT`	Phase 1 commit
`P2-COMIT`	Phase 2 commit
`SIGN OFF`	User signoff
`SIGN ON`	User signon
`SQL-xxxx-UNK`	Unrecognized SQL call code; `xxxx` is the hexadecimal code value
`T_IDENT`	Terminate identify
`T_THREAD`	Terminate thread

SQL DB2 Calls

The format DTRAC uses for SQL DB2 calls is shown below.

Format:

- DB2 subsystem identifier
- SQL statement call type

Example:

```
DB2S  SELECT
```

SQL statement call types:

DTRAC calls are shown to the left. When a DTRAC call has been abbreviated from the complete form of a call, the complete form is shown in parentheses to the right of the abbreviation.

```
ALTER DB  (ALTER DATABASE)        DROP PK   (DROP PACKAGESET)
ALTER IX  (ALTER INDEX)           DROP SG   (DROP STORAGEGROUP)
ALTER SG  (ALTER STORAGEGROUP)    DROP SY   (DROP SYNONYM)
ALTER SY  (ALTER SYNONYM)         DROP TB   (DROP TABLE)
ALTER TB  (ALTER TABLE)           DROP TS   (DROP TABLESPACE)
ALTER TS  (ALTER TABLESPACE)      DROP VW   (DROP VIEW)
CLOSE                             EXECI     (EXECUTE IMMEDIATE)
COMMENT   (COMMENT ON)            EXECUTE
CONN TO   (CONNECT TO)            EXPLAIN
CONN RST  (CONNECT RESET)         FETCH
CONNECT                           GRANT
CREAT AL  (CREATE ALIAS)          IMPL CON  (IMPLICIT CONNECT)
CREAT DB  (CREATE DATABASE)       INSERT
CREAT IX  (CREATE INDEX)          LABEL     (LABEL ON)
CREAT SG  (CREATE STORAGEGROUP)   LOCK
CREAT SY  (CREATE SYNONYM)        OPEN
CREAT TB  (CREATE TABLE)          PREPARE
CREAT TS  (CREATE TABLESPACE)     REVOKE
CREAT VW  (CREATE VIEW)           RMT SQL   (REMOTE SQL)
DELETE                            SELECT
DESCRIBE                          SET C SQ  (SET CURRENT SQLID)
DROP AL   (DROP ALIAS)            SET C PK  (SET CURRENT PACKAGESET)
DROP DB   (DROP DATABASE)         SET H VA  (SET HOST VARIABLE)
DROP IX   (DROP INDEX)            UPDATE
```

MQSeries Calls

This section describes the formatting used in the DTRAC EVENT column for events resulting from MQSeries calls.

Format:

- MQSeries subsystem identifier
- MQSeries call type or IMS attach facility call type

Examples:

```
CSQ3 C_THREAD
CSQ3 MQCONN
CSQ3 MQPUT1
```

MQSeries call types:

MQBACK	Backs out updates to queue manager resources
MQCLOSE	Closes a message queue
MQCMT	Commits updates to queue manager resources
MQCONN	Sets up contact with a queue manager
MQDISC	Terminates contact with a queue manager
MQGET	Takes messages from a queue
MQINQ	Requests information about a message queue
MQOPEN	Opens a message queue
MQPUT	Puts messages on a queue
MQPUT1	Puts a single message on a queue
MQSET	Sets attributes on a queue

IMS attach facility call types:

xx-UNKC	Unrecognized call type; xx is the hex value of the call type.
ABORT	Abort call
C_THREAD	Create thread
IDENTIFY	Identify
INIT_CAL	Initiate call
P1-COMIT	Phase 1 commit
P2-COMIT	Phase 2 commit
SIGN OFF	User signoff
SIGN ON	User signon
T_IDENT	Terminate identify
T_THREAD	Terminate thread

Part 5. Monitor Service Control

This section describes how to control time-driven monitor services using the SET timer request and the BBI Subsystem Information application (Primary Menu Option 5) from the Primary Option Menu.

Chapter 37. SET Timer Request

The SET request is used to initiate, modify, or purge timer-driven services for:

- Data collection monitor, workload wait, or workload trace services.
- Service output display Image logging. It is activated from the service lists. When a service is selected, a data entry panel is displayed and primed with SET request options for the selected service.

Service Requests

SET is used to request timer-driven:

- Data collection by resource, workload wait, or workload trace monitor services
- Image logging of IMS analyzer or monitor status displays

SET options for a service are primed in a data entry panel when that service is requested. SET for a service also can be requested from the SERV field of a display or as a member of a BBPARM data set to initiate multiple service requests.

Request Initiation

A SET request for a service starts with one of the following keywords:

`REQ=reqid` Defines a new request or continues a held request where `reqid` is the service select code for a service and an optional parameter identifying the request.

`MOD=reqid` Modifies an existing request.

`BLK=membername` Identifies a BBPARM data set member that contains a group of service requests.

Service Request ID (reqid)

A service is uniquely defined in a SET request by its `reqid`. The `reqid` is the service select code (ssc) and, if necessary, a parameter:

```
REQ=ssc,parameter
```

Duplicate requests are not allowed; however, multiple requests for the same service can be active concurrently if the `reqid` for each request has a unique parameter.

The original `reqid` must be specified to modify or purge an active service request.

Syntax

If you specify a SET request from the SERV field of a display panel or in a member of a BBPARM data set as described in the following sections, the request requires a syntax.

Note: If you select a service from a list application, there are no syntax requirements. A data entry panel is displayed that is primed with the appropriate keywords for the selected service.

The syntax for entering keyword options (`REQ=`) in the PARM field of a display panel or a BBPARM data set is free format and keyword-oriented. Any number of blanks, commas, slashes, or parentheses can be interspersed in the text between keywords to improve readability; they are ignored during request processing. The equal sign between keyword and operand is optional, but it is recommended to improve readability.

Single Requests

Single SET request keyword options for a service can be entered in the PARM field of a service display panel. The PARM field is 55 bytes long. As many keyword options as can fit in this space can be entered. If all the options do not fit in the parameter line, the option `HOLD=YES` can be used to hold the request until it can be completed with a second SET request with the same `reqid`.

Multiple Requests

A series of valid request keyword options for SET can be predefined in a member of the BBI-SS PAS BBPARM data set. In the BBPARM data set there are sample block request members to set and purge multiple requests. BLKDBCW is a suggested starter set of monitors. BLKDBCWP is a sample member to purge the requests set by BLKDBCW.

Each request can start anywhere in a line from columns 1 to 79 and can continue onto the next line.

There can be one group for each IMS subsystem. Requests for display logging and starting monitors can be specified; for example, member BLKIMSA in BBPARM could contain:

```
REQ=PI I=00:06:00 LOG=ATINTVL
REQ=FPBST   I=00:01:00
REQ=DWARN   I=00:01:00
REQ=DBHIT,WVAL=80,WLIM=3,LOG=ATINTVL
REQ=DBTOT,WVAL=60,WLIM=3,LOG=ATINTVL
REQ=LHELD,WVAL=90,WMSG=WTO,WIF=2,WIN=5,WLIM=6,LOG=ATINTVL
```

The member can be started by specifying the BLK keyword and BBPARM member name in the parameter field for the SET timer request:

```
SERV ==> SET
PARM ==> BLK=BLKIMSA
```

Note: The BLK keyword cannot be used in a request within a BBPARM member.

Multiple BBPARM members with predefined requests can be included with the BLK request:

```
BLK=membername1(,membername2)(,...)
```

The member names must identify previously created members of a partitioned data set, like BBPARM, allocated to the BBI-SS PAS with the ddname BBIPARM (see the SSJCL sample member in BBSAMP). The record length must be 80 characters; blocksize can be any multiple of 80.

A SET BLK request is logged to the BBI-SS PAS Journal log. Any errors in the requests issued with SET BLK also are logged to the BBI-SS PAS Journal log.

Automatic Startup of Multiple Requests

Multiple requests can be started when the BBI-SS PAS starts. The name of the BBPARM member containing the block of requests is specified in the BBIISP00 member of BBPARM. The parameters are TARGET, BLK, USRID, and AUTOID. They are used as follows:

TARGET — Defines the IMS subsystem to be monitored. It must be on the same line as BLK and USRID.

BLK — Specifies a BBPARM member that contains a block of monitor requests to be started for the specified target IMS subsystem when the BBI-SS PAS starts. It must be on the same line as TARGET and USRID.

USRID — For each specified target IMS subsystem, USRID is the security authorization ID for each monitor request within the BBPARM member specified by BLK. It must be on the same line as TARGET and BLK. USRID overrides AUTOID.

AUTOID — Is a security authorization ID that applies to all the specified target IMS subsystems unless USRID is specified. AUTOID is written on a separate line from TARGET and BLK.

TARGET, BLK, and USRID must be written on the same line and can be repeated for different targets. TARGET and BLK are required; USRID and AUTOID are optional. If you do not specify USRID or AUTOID, BBPARM member USERID is used for the security authorization ID. If this is not a valid authorization ID, BMC Software strongly recommends that you specify either USRID or AUTOID.

For example:

```
TARGET=IMSA,BLK=BLKIMSA,USRID=$TIMSAR
TARGET=IMSB,BLK=BLKIMSB,USRID=$TIMSBX
TARGET=IMSC,BLK=BLKIMSC,USRID=$TIMSC
```

or

```
AUTOID=ADMIN
TARGET=IMSA,BLK=BLKIMSA
TARGET=IMSB,BLK=BLKIMSB
TARGET=IMSC,BLK=BLKIMSC
```

The first example assumes there are three active IMS subsystems. The monitors defined in BLKIMSA extract data from the IMSA subsystem and are associated with the user ID $TIMSAR. The monitors defined in BLKIMSB extract data from the IMSB subsystem and are associated with the user ID $TIMSBX.

Note: If the target IMS is not active, the QIS option defines the action to be taken for each request. The default is to quiesce until the target IMS starts.

Multiple Request Comments

Comments can be entered in a block request member of BBPARM. A comment is delimited by the * sign. A warning indicator is issued when invalid characters, which are treated as blanks, are encountered.

Note: If line numbers are used in a block request member, each number should be preceded by an * sign.

Request Termination

A service request can be terminated with the following SET request:

```
SERV ==> SET
PARM ==> PRG=reqid              Purges an existing request
```

SET Keyword Parameter Options

Table 15 on page 305 and Table 16 on page 308 describe the keyword parameters that can be specified with a SET request.

Nonmodifiable Keyword Options

The following options cannot be modified (`MOD=reqid`) because previously collected history would be distorted.

INTERVAL
RANGES
START
TITLE
HOLD

HOLD cannot be used for a request being modified. If a change needs to be made to one of these options for an active request, the request must be purged and a new request entered.

Selection Criteria

Workload requests can be qualified with selection criteria keywords as described in Table 9 on page 210. Once a workload request is activated, the selection criteria keywords cannot be changed. The request must be stopped and reactivated.

A + character can be used as a generic name qualifier for a workload request.

Keywords

The SET keywords define a timer request function. Table 15 lists the keywords by function. They are used to:

- Make a SET request from a display panel
- Make a SET request from within a BBPARM data set member
- Prime the data entry panels used to make a request for a time-driven data collection service.

Table 15 is an index to the keywords used to make a SET request for a service. These keywords are described in:

- “Using the Resource Monitor Data Entry Panel” on page 95
- “Using the Workload Monitor Data Entry Panel” on page 100
- “Using the Workload Global Region Call Monitor Data Entry Panel” on page 107
- “Using the Workload Trace Data Entry Panel” on page 205
- “Using the Workload Wait Data Entry Panel” in the MWAIT chapter in *MAINVIEW for IMS Online – Analyzers Reference Manual*

Any of the keywords can be used to make block requests as described in “Multiple Requests” on page 302. Table 16 on page 308 describes the keywords that affect all service activation, including multiple requests defined in a BBPARM member or from the SERV field of a display panel.

Table 15. SET Keywords

Options	**Keywords**
SET Keywords to Define Requests	BLK MOD PRG REQ
SET Keywords to Define Request Activation	HOLD INTERVAL RST START STOP STOPCNT TARGET|TGT
SET Keywords to Define Warnings (monitors only)	WIF WIN WLIM WMAX|WVAL WMSG

Table 15. SET Keywords (continued)

Options	Keywords
SET Keywords to Define Special Options	LOG PLOTMAX QIS RANGES TITLE|T
SET Keywords to Define Workload Analyzer Parameters for MWAIT and MTRAC	CURPER HISTORY STORAGE TYPE WRAP
SET Keywords to Define Selection Criteria for Workload Analyzer MWAIT and MTRAC Services and Workload Monitor Services AND logic is used for keywords. OR logic is used for keyword operands.	ABEND CALLTYPE|CTYP CLASS DATABASE DB2PLAN TERM USERID PGMTYPE|PTYP PROG PSB REGION RGNID TRAN TRANTYPE|TTYP
SET Keywords to Define Trace Logging for MTRAC Service	TRARCSTC TRCYL TRDISP TRDSN TRNUMDS TRREUSE TRSMSDCL TRSMSMCL TRSMSSCL TRSWTIME TRSUFFIX TRVOLS

Table 15. SET Keywords (continued)

Options	Keywords
SET Keywords to Define Exception Filters for MTRAC Service	DB2CPU DBUSED DEPCPU DLICALL DLIDC DLIDLET DLIGTYP DLIISRT DLIMSGN DLIMSIN DLIMSPR DLIREPL DLITOT DLIUTYP ELAPSED FPCICON FPOBA FPWTBFR INPUTQ RDWRT RESPONS SERIO SQLDDL SQLDEL SQLDYN SQLINS SQLOPN SQLOTH SQLSEL SQLTOT SQLUPD TOTCPU

Table 16 describes the keywords that affect all service activation, including multiple requests defined in a BBPARM member or from the SERV field of a display panel

Table 16. SET Keywords Affecting Request Activation

Keyword	Operand	Description
BLK		Identifies a member in BBPARM that contains predefined service requests.
	name	Member name in BBPARM.
	RRR	Resets all requests according to their RST options. This is the same as starting IMS after the requests were quiesced. It could be used at midnight to reset statistics for daily monitors.
	SSS	Restart all quiesced requests according to their RST options after IMS is restarted.
	ZZZ	Internal use only.
HOLD	YES	Holds this request until the definition is completed with another SET with the same reqid. Use it to specify more keywords than will fit on the 55-byte parameter line of a service display. **Note:** The MOD keyword cannot be used to change this option. The request must be purged and a new request must be made.
	NO	Default in NO.
INTERVAL|I	hh:mm:ss	The time interval between successive invocations of the requested service. The default is one minute (00:01:00) or as specified by the user in the BBIISP00 member of the BBPARM data set. It can be used with the LOG keyword to request automatic logging of a display to the BBI-SS PAS Image log. The MOD keyword cannot be used to change this option. The request must be purged and a new request must be made.

Table 16. SET Keywords Affecting Request Activation (continued)

Keyword	Operand	Description
LOG		Specifies if and when automatic logging occurs. Analyzer, DMON or DWARN, or PLOT displays can be logged to the BBI-SS PAS Image log. For the monitors, a PLOT of the latest data is written to the BBI-SS PAS Image log.
	NO	No logging. Default for monitor services. `LOG=NO` is the only valid option for monitors that only measure a condition against a warning threshold (data measurement type of warning only).
	ATSTOP	Logs display when processing of this request is stopped. If `QIS=Y` has been specified in the request, `LOG=ATSTOP` is invoked at IMS termination and at BBI-SS PAS termination.
	ATPD	Logs display at each period of 10 intervals.
	ATINTVL	Logs display at each interval as specified by the user with the INTERVAL parameter or in the BBIISP00 member of the BBPARM data set. Default for analyzers. ATINTVL is the default for logging images of the analyzer services to the BBI-SS PAS Image log.
	ATWARN	Logs a plot whenever a warning message is generated by the associated monitor. NO is the default for the monitor services.
MOD	reqid	Modifies an existing request.
PARM	`id`	Where `id` is a unique 1- to 8-character identifier. If you want to run multiple requests, specify an ID in the PARM field to make this request unique. Blank can also be used as an identifier for one request. The identifier is used to specify which request is to be logged to the BBI-SS PAS Image log for an image log request, to be displayed with the plot display service for a monitor request, or to be displayed by a trace display service for a trace request.
PRG		Purges existing request or requests (frees the request block).
	reqid	Purges a single request.
	ALL	Purges all existing requests if the user is authorized.
REQ	reqid	Defines a new request or continues a held request.

Table 16. SET Keywords Affecting Request Activation (continued)

Keyword	Operand	Description
QIS		Defines the action to be taken for the requested service when IMS is not active.
	YES	Specifies that the requested service is to be quiesced. This is the default.
	NO	Specifies that the requested service is to start or continue running. When `QIS=NO` is specified, monitors that require IMS continue to be scheduled at each interval; however, they return zero values. Any analyzer services set up to run asynchronously fail with a short message of CANNOT LOCATE IMS SPECIFIED in the first line. The BBI-SS PAS Image log contains screen images of these services.
RST		Defines the restart option to be used when a service or Image log request is quiesced because of an inactive IMS subsystem or `RRR` is specified for BLK (see “Multiple Requests” on page 302). Default is HOT.
	HOT	Restarts the service request automatically without loss of previous data.
	COLD	Restarts the service request automatically; all previously collected data is deleted.
	PUR	Purges the service request automatically when the target IMS starts.
	QIS	Keeps the service in a quiesced state until it is purged by an authorized user.
START	hh:mm:ss	Requests a start time for the service request. If the time entered is more than 10 minutes prior to the current time, 24 hours are added to the specified time and the request is started the next day. To start a request at midnight, specify 24:00:00. The default is the next full minute. **Note:** This option cannot be modified. The request must be purged and a new request must be made.

Table 16. SET Keywords Affecting Request Activation (continued)

Keyword	**Operand**	**Description**
STOP	`nn|hh:mm:ss`	Requests a stop limit for a service request where `nn` is length of time in minutes and `hh:mm:ss` is a timestamp. Processing stops at the end of the last interval before the specified stop time. This time is displayed in the STOP field when the request is viewed with the R, P, M, or W line commands from the Active Timer Requests list (Primary Option Menu 2) application. If the time entered is the same as the START time, 24 hours are added to the STOP time. **Note:** For block requests, you must specify either `hh:mm:ss` or `STOPCNT=nn`
STOPCNT	n	The number of intervals to be processed. Requests that have completed without collecting history data are purged. Default is no limit. Requests are processed until stopped or purged or until the BBI-SS PAS terminates. **Note:** For block requests, you must specify either `STOP=hh:mm:ss` or `STOPCNT=nn` where `nn` is the number of minutes until the monitor stops.
TARGET| TGT	job name	TARGET is used to override the default target IMS identified in the TGT field. The new target IMS will be that specified by the name parameter. Valid entries are a 1- to 8-character IMS control region job name or step name. IMSID can be used if identified in BBPARM member BBIJNT00. MAINVIEW AutoOPERATOR must be installed. TARGET is required if the IMFC command is used in a MAINVIEW AutoOPERATOR EXEC to request analyzer or monitor services. A message is issued to the BBI-SS PAS Journal log if this keyword is not coded in the request.

Chapter 38. Displaying BBI-SS PAS Information (Primary Menu Option 5)

This application is used to control all requests for timer-driven services, which include:

- Data collection by the Resource Monitor, Workload Monitor, and Workload Analyzer wait and trace monitor services
- Image logging of the analyzer or monitor service displays

Selecting Primary Menu Option 5 from the Primary Option Menu displays general information about the Timer facility, as shown by Figure 26.

```
BMC SOFTWARE  ---------- DISPLAY STATISTICS AND DEFAULTS ----- PERFORMANCE MGMT
COMMAND ===>                                                TGT ===> IMSxxx
                                                            TIME --  14:36:03

BBI-SSID: RS33         BBI Release -- 2.6.0      SS Started: 10:51:25 22-MAR-02
MVIMS/EC 3.3.10        MVIMS/PAS 3.3.10          SS Elapsed: 03:44:37
EC TRACE BUFFERS= 0    IN USE= 0    SIZE= OK     PUT LEVEL= 0201C
MVIMS/IMS Components WA: D WM: D PR: D TA: D
MSIMS/SS  Components WA: D WM: D RA: Y RM: Y LM: Y
   Requests:           80 Total request blocks        56 Unused blocks
   Activity:       39,090 Service calls               20 Warnings written

   Parameters:  ---- General ----     -- IMS WA WAIT --     - TRACE BUFFERS -
                INTERVAL=00:01:00     CURPER=00:01:00       STORAGE=100K
                ROUT=NONE             HISTORY=00:30:00      TRBUFF=10
                DESC=NONE                                   TRSIZE=4K

------------------------ DEFINED REQUESTS BY TARGET --------------------------
LC  TARGET   TYPE     ACTIVE  INIT  COMPLETE  HELD  INVALID  LOCKED   QIS   RST
    -TOTAL-  --ALL--      24
    IMSxxx   AO-CALLX      1
    IMSxxx   MONITOR      23
     **************************** END OF REQUESTS *****************************
```

Figure 26. Sample BBI Information Display (Primary Menu Option 5)

This application shows the BBI-SS PAS status; status of the currently active Event Collector, analyzer, monitor, and offline components; timer facility activity statistics; and timer request default parameters in effect and summarizes the status of all the timer requests.
The information shown is for the BBI-SS PAS connected to the target specified in the TGT field of the display.

The panel components are:

BBI-SS PAS Status Information
Component Status Information
Timer Facility Activity Statistics
Timer Request Default Parameters
Defined Timer Requests by Target

The panel fields are shown and described by panel component in the sections that follow.

BBI-SS PAS Status Information

This area displays BBI-SS PAS status information.

```
BBI-SSID: RS33        BBI Release -- 2.6.0      SS Started: 10:51:25 22-MAR-02
                                                SS Elapsed: 03:44:37
                                                PUT LEVEL= 0201C
```

Figure 27. BBI-SS PAS Status (Timer Statistics and Defaults)

These fields show the status of the BBI-SS PAS by:

Field	Description
BBI-SSID	A two- to four-character code that identifies the active SS (subsystem).
BBI Release	The installed release level (not the product release) of the BBI platform.
SS Started	The date and time the BBI-SS PAS was started.
SS Elapsed	The length of time the BBI-SS PAS has been active.
PUT LEVEL	The level of the latest installed PUT tape

Component Status Information

This area displays the status of the currently active Event Collector, Workload Analyzer, Workload Monitor, Resource Analyzer, Resource Monitor, Performance Reporter, and Transaction Accountant.

```
MVIMS/EC 3.3.10        MVIMS/PAS 3.3.10
EC TRACE BUFFERS= 0    IN USE= 0    SIZE= OK
MVIMS/IMS Components WA: D WM: D PR: D TA: D
MSIMS/SS  Components WA: D WM: D RA: Y RM: Y LM: Y
```

Figure 28. Product and Component Status (Timer Statistics and Defaults)

These fields show the status of the currently active components.

Field	Description
MVIMS/EC	Active release of the Event Collector.
MVIMS/PAS	Active release of MVIMS or MVDBC.
EC TRACE BUFFERS	Number of allocated Event Collector trace buffers.
IN USE	Number of allocated Event Collector trace buffers actually being used.

SIZE	Actual size of the Event Collector trace buffers.

Note: If you change TRBUFF and TRSIZE in BBPARM member IMFBEX00 and use the following RESET BBI-2 control command to reset them to your changed values:

```
.RESET PARM IMFBEX00
```

your changes do not take effect until the IN USE field in this display is zero.

For a complete description of the RESET command, see the *MAINVIEW Common Customization Guide*.

MVIMS/IMS Components	Product components active in the IMS control region. N indicates not active, Y indicates active, and D indicates active with DB2 extensions. The components are shown as: WA Workload Analyzer WM Workload Monitor PR Performance Reporter TA Transaction Accountant
MVIMS/SS Components	Components active in the BBI-SS PAS. N indicates not active, Y indicates active, and D indicates active with DB2 extensions. The components are shown as: WA Workload Analyzer WM Workload Monitor RA Resource Analyzer RM Resource Monitor LM Resource Analyzer and Resource Monitor

Timer Facility Activity Statistics

This area displays timer facility activity statistics.

```
Requests:          80 Total request blocks        56 Unused blocks
Activity:      39,090 Service calls               20 Warnings written
```

Figure 29. Timer Facility Activity (Timer Statistics and Defaults)

These fields show the timer facility activity by:

Field	**Description**
Requests	Total Request Blocks The maximum number of timer requests that can be defined concurrently as specified with the MAXREQ parameter in BBPARM member BBIISP00 (see “BBI Libraries and Data Sets” in the *MAINVIEW Common Customization Guide* for a description of BBIISP00). Unused Blocks The number of blocks that are still available for new timer requests.
Activity	Service Calls The number of times the timer facility has invoked a service. This number includes requests for data collection monitors and automatic Image logging of analyzer or monitor displays. Warnings Written The number of initial warning condition messages issued by the requested data collection monitors.

Active Default Parameters

This area displays active default parameters.

```
  Parameters:  ---- General ----     -- IMS WA WAIT --     - TRACE BUFFERS -
               INTERVAL=00:01:00     CURPER=00:01:00       STORAGE=100K
               ROUT=NONE             HISTORY=00:30:00      TRBUFF=10
               DESC=NONE                                   TRSIZE=4K
```

Figure 30. Active Default Parameters (Timer Statistics and Defaults)

These fields show the timer facility default parameters in effect. The defaults are defined in the BBIISP00 member of the BBPARM data set. These are described below in alphabetical order.

Field	Description
CURPER	The interval specified for the CURRENT PERIOD area of the DWAIT (workload wait data) display.
DESC	The descriptor code(s) for monitor warning WTO messages. NONE is the default.
HISTORY	The interval specified for the HISTORY area of the DWAIT (workload wait data) display.
INTERVAL	The default timer request interval specified in the BBITSP00 member of the BBPROF data set. This value is used if the INTERVAL keyword is not specified when a timer request is defined.
ROUT	The OS/390 console route code(s) for monitor warning WTO messages. NONE is the default.
STORAGE	The amount of extended BBI-SS PAS private area storage allocated for the trace entry buffer.
TRBUFF	The number of trace buffers allocated for each active detail trace.
TRSIZE	The size of each trace buffer.

The IBM messages and codes manual describes the codes that may appear in the ROUT and DESC fields.

Defined Requests by Target

This area displays defined requests by target.

```
------------------------ DEFINED REQUESTS BY TARGET ---------------------------
LC  TARGET   TYPE     ACTIVE  INIT  COMPLETE  HELD  INVALID  LOCKED   QIS   RST
    -TOTAL-  --ALL--      24
    IMSxxx   AO-CALLX      1
    IMSxxx   MONITOR      23
```

Figure 31. Request Summary by Target. Timer Statistics and Defaults

This portion of the Timer Statistics panel is a scrollable list of all the requests per target for the BBI-SS PAS shown in the BBI-SS PASID field (see “BBI-SS PAS Status Information” on page 314). It shows the request type for each target and the amount of activity for each request state by:

Field	**Description**
LC	A line command input field. A one-character line command (see “Line Commands” on page 319) can be entered in this field to select a list display of the active timer requests, as described in “Active Timer Requests (S Line Command)” on page 319.
TARGET	IMS job name or ID.
TYPE	Types of timer requests:

MONITOR	Monitor service requests
IMG-LOG	Automatic Image logging requests of analyzer or monitor displays
AO-CALLX	Time-initiated EXEC requests that use the CALLX service.

The total number of requests per target is shown for each of the following request states:

ACTIVE	Active requests.
INIT	Requests waiting to be invoked (a start time was specified, but it has not been reached).
COMPLETE	Requests that completed normal execution.
HELD	Requests being held and pending release.
INVALID	Requests that terminated because of an invalid parameter or measurement. The BBI-SS PAS Journal log contains descriptive messages of the request errors.
LOCKED	Requests that terminated because of a LOCK command or a service routine ABEND.
QIS	Requests that quiesced because the target IMS subsystem was not active.

RST The number of requests waiting until the current interval expires before performing restart processing after the target IMS subsystem has restarted. It is specified by the RST keyword in the original request.

Line Commands

Entering the following line command in the LC field of the Timer Facility application executes the line command function:

Line Command	Description
S	SELECT. Selects the Active Timer Requests application showing: • All the BBI-SS PAS requests (see Figure 32). S is entered in the LC field for the TOTAL targets (see Figure 31 on page 318). • Only those requests for a specific target. S is entered in the LC field for the target identifier (see Figure 31 on page 318).

Active Timer Requests (S Line Command)

The S line command displays the Active Timer Request application, which is described in Chapter 6, “Displaying a List of Active Timer Requests (Primary Menu Option 2)” on page 65. The list shown in Figure 32 is displayed when the S line command is entered in the LC input field for TOTAL. It displays all the timer requests active for the target shown in the TGT field.

```
BMC SOFTWARE  -------------- ACTIVE TIMER REQUESTS --------- COMMAND(S) ISSUED
COMMAND ===>                                                  TGT ===> IMSxxx

                                    INPUT  INTVL ==> 3        TIME --  14:32:27
  COMMANDS: SM (START MONITORS), SORT, AREA, X ON/OFF, DM (DMON), DW (DWARN)
  LC CMDS:  S (SELECT), W (SHOW),      M (MODIFY)
            P (PURGE),  R (REPLICATE), H (HELP), Z (STOP)                    >>>
  LC   SERV  PARM     TITLE                 CURRENT  WVAL  -8-6-4-2-0+2+4+6+8+
       DBIO           DB I/O COUNT BY SUBPOOL     36    10| *********W********* |
       DBIO  1        DB I/O COUNT BY SUBPOOL      0      |                     |
       DBIO  3        DB I/O COUNT BY SUBPOOL     13      |                     |
       DBIO  2        DB I/O COUNT BY SUBPOOL      0      |                     |
       DBHIT          OSAM HIT RATIO BY SP        93   <80| *******  W          |
       DBHIT 1        OSAM HIT RATIO BY SP         0   <90|          W          |
       DBHIT 3        OSAM HIT RATIO BY SP       100   <80| ******   W          |
       VHIT           VSAM HIT RATIO BY SP       100      |                     |
       OBUFW          OLDS BUFFER WAITS            0     1|          W          |
       OCHKW          OLDS CHECK WRITES           13     1| *********W********* |
       WADIO          WADS I/O                    15     1| *********W********* |
       DPAGE          DEMAND PAGING BY REGION    498   500| *********W          |
       CSAUT          CSA % UTILIZATION           56    80| ******   W          |
       ECSAU          ECSA % UTILIZATION          61    60| *********W          |
       DLIO  ++++++++ DL/1 EXCP COUNT BY DDNAM    31      |                     |
       SYSIO ++++++++ EXCP COUNT BY DDNAME        21      |                     |
       SHMSG          SHORT MSG QUEUE % UTIL       0    10|          W          |
       LGMSG          LONG MSG QUEUE % UTIL        4    10| ***      W          |
       QBLKS          QBLKS % UTILIZATION          0    10|          W          |
       LHELD          IRLM LOCKS HELD              2      |                     |
       LKREQ          IRLM LOCK REQUESTS           0      |                     |
       LSUSP          IRLM SUSPENSIONS             0      |                     |
```

Figure 32. Active Timer Requests List (Timer Facility Application S Line Command)

Part 6. Appendixes

Appendix A. IMS Dump Analysis

This appendix describes how to analyze an IMS dump with MAINVIEW for IMS or MAINVIEW AutoOPERATOR for IMS installed.

AO Exit

Except for MAINVIEW AutoOPERATOR for IMS MTO message capture for the Journal log, all AO code merely passes control to the specified routines during operation. During initialization, special protection exists while the Event Collector is being set up.

MAINVIEW AutoOPERATOR for IMS Routines in IMS

During initialization, the MAINVIEW AutoOPERATOR for IMS AO code creates two subtasks under the IMS control task. Each of these is protected by ESTAE routines and uses different control blocks than IMS. These subtasks generally can be ignored during IMS dump analysis since they do not affect the IMS flow. These subtasks are terminated correctly at IMS termination.

Event Collector

The following information should be noted about IMS dumps:

- MAINVIEW for IMS register save areas are in MAINVIEW for IMS data areas, not in the IMS prechained save areas. The IMS chains remain unchanged.
- Sometimes the R14 return register in an IMS save area does not point back into the calling IMS module. The IMS R14 value can be found 4 bytes in front of the address pointed to by R14.
- A MAINVIEW for IMS module at entry saves the registers of an IMS module in the next IMS prechained save area pointed to by R13.
- MAINVIEW for IMS module registers are always saved in MAINVIEW for IMS save areas. These save areas are assigned dynamically as required. Normally, one of several preallocated save areas per region is used.

  Each active IMS region has a MAINVIEW for IMS data area acquired for it at region initialization. This data area is in ECSA and is named IMERD*nnn*, where *nnn* is the PST number. The preallocated save areas are in this block.

  If more save areas are required, a dynamic storage pool is used. This pool is also in ECSA and is named IMFSP000.

  Any area in actual use as a MAINVIEW for IMS save area, whether in IMERDnnn or IMFSP000, is identified with SAR or ISA.
- In most cases, only a save area backward pointer exists, pointing from the MAINVIEW for IMS area to the previous IMS save area. While a MAINVIEW for IMS module is in control, the current R12 is its base register and the current R13 points to its SAR.

- In some cases, a MAINVIEW for IMS module transfers control to an IMS module, but needs to regain control after it completes processing. In this case also, the IMS save area chains remain unchanged.

  The only sign of the presence of a MAINVIEW for IMS module is an R14 value in an IMS save area which does not point back to the calling IMS module. If you need to verify the IMS path, the IMS R14 value is saved in the MAINVIEW for IMS SAR, 4 bytes in front of the address pointed to by the R14 in the IMS save area.

  Note: When MAINVIEW for IMS interfaces between two IMS modules, MAINVIEW for IMS is transparent to the IMS modules. All registers are preserved.

- IMECSRVx and IMFCSRVx appear in dumps as active ITASKS. This is normal as long as the current save area is for DFSIWAIT.

Appendix B. How Product Libraries Should Be Used

Several distributed libraries are included with your MAINVIEW products, including a parameter library (BBPARM), a sample library (BBSAMP), and a profile library (BBPROF). Use the contents of these distributed libraries as models to create site-customized product libraries, either manually or automatically, with AutoCustomization.

Warning

The distributed libraries should never be modified. If you change the distributed libraries, subsequent SMP maintenance will overwrite your changes.

Throughout the MAINVIEW documentation set, references to these libraries use the distributed name. However, when you need to make changes, be sure to use the corresponding library that has been customized for your site. Table 17 lists the distributed name, the corresponding customized library created by AutoCustomization, and leaves space for you to note any other corresponding library that may have been created for your site.

Table 17. Product Libraries

Distributed Library Name	**Library Created by AutoCustomization**	**Other Site-Customized Copy**
BBPARM	UBBPARM	
BBSAMP	UBBSAMP	
BBPROF	SBBPROF	

For more detailed information about all the product libraries, see “Using MAINVIEW Product Libraries” in the *MAINVIEW Common Customization Guide* or “Using Product Libraries” in the *MAINVIEW Administration Guide*.

Index

Symbols

A

B

C

E

F

G

N

O

P

Q

R

S

T

Notes

41593